Revocable Cryptography from Learning with Errors

Prabhanjan Ananth* UCSB
Alexander Poremba† Caltech
Vinod Vaikuntanathan‡ MIT

Abstract

Quantum cryptography leverages many unique features of quantum information in order to construct cryptographic primitives that are oftentimes impossible classically. In this work, we build on the no-cloning principle of quantum mechanics and design cryptographic schemes with *key-revocation capabilities*. We consider schemes where secret keys are represented as quantum states with the guarantee that, once the secret key is successfully revoked from a user, they no longer have the ability to perform the same functionality as before.

We define and construct several fundamental cryptographic primitives with *key-revocation capabilities*, namely pseudorandom functions, secret-key and public-key encryption, and even fully homomorphic encryption. Our constructions either assume the quantum subexponential hardness of the learning with errors problem or are based on new conjectures. Central to all our constructions is our approach for making the Dual-Regev encryption scheme (Gentry, Peikert and Vaikuntanathan, STOC 2008) revocable.

*prabhanjan@cs.ucsb.edu
†aporemba@caltech.edu
‡vinodv@mit.edu

Contents

1 Introduction

Quantum computing presents exciting new opportunities for cryptography, using remarkable properties of quantum information to construct cryptographic primitives that are unattainable classically. At the heart of quantum cryptography lies the *no-cloning principle* [WZ82, Die82] of quantum information which stipulates that it is fundamentally impossible to copy an unknown quantum state. Indeed, Wiesner [Wie83] in his seminal work from the 1970s used the no-cloning principle to construct a quantum money scheme, wherein quantum states are used to construct banknotes that can be verified to be authentic (using a secret key) but cannot be counterfeited. Ever since this watershed moment, and especially so in the recent years, a wide variety of primitives referred to as *unclonable* primitives have been studied and constructed in the context of encryption [Got02, BL20, BI20b, GZ20], digital signatures [LLQZ22] and pseudorandom functions [CLLZ21].

Our Work: Revocable Cryptography. Delegation and revocation of privilege are problems of great importance in cryptography. Indeed, the problem of revocation in the context of digital signatures and certificates in the classical world is a thorny problem [Stu95, Riv98]. In this work, we undertake a systematic study of *revocable (quantum) cryptography* which allows us to delegate and revoke privileges in the context of several fundamental cryptographic primitives. This continues a recent line of work in quantum cryptography dealing with revoking (or certifiably deleting) states such as quantum ciphertexts or simple quantum programs [Unr13, BI20b, GZ20, AL21, HMNY21a, Por22, BK22]. In this framework, revocation is to be understood from the perspective of the recipient of the quantum state; namely, the recipient can certify the loss of certain privileges by producing a certificate (either classical or quantum) which can be verified by another party.

As a motivating example, consider the setting of an employee at a company who takes a vacation and wishes to authorize a colleague to perform certain tasks on her behalf, tasks that involve handling sensitive data. Since the sensitive data is (required to be) encrypted, the employee must necessarily share her decryption keys with her colleague. When she returns from vacation, she would like to have her decryption key back; naturally, one would like to ensure that her colleague should not be able to decrypt future ciphertexts (which are encrypted under the same public key) once the key is "returned". Evidently, if the decryption key is a classical object, this is impossible to achieve as classical keys can be copied at will.

In revocable (quantum) cryptography, we associate a cryptographic functionality, such as decryption using a secret key, with a quantum state in such a way that a user can compute this functionality if and only if they are in possession of the quantum state. We then design a revocation algorithm which enables the user to certifiably return the quantum state to the owner. Security requires that once the user returns the state (via our revocation algorithm), they should not have the ability to evaluate the functionality (e.g. decrypt ciphertexts) anymore. We refer to this new security notion as *revocation security*.

Another, possibly non-obvious, application is to detecting malware attacks. Consider a malicious party who hacks into an electronic device and manages to steal a user's decryption keys. If cryptographic keys are represented by classical bits, it is inherently challenging to detect such attacks that compromise user keys. For all we know, the intruder could have stolen the user's decryption keys without leaving a trace. Indeed, a few years ago, decryption keys which were used to protect cell-phone communications [Int15] were successfully stolen by spies without being

detected.[1] With revocable cryptography, a malicious user successfully stealing a user key would invariably revoke the decryption capability from the user. This latter event can be detected.

Our Results in a Nutshell. We construct revocable cryptographic objects under standard cryptographic assumptions. Our first main result constructs a key-revocable public-key encryption scheme, and our second main result constructs a key-revocable pseudorandom function. We obtain several corollaries and extensions, including key-revocable secret-key encryption and key-revocable fully homomorphic encryption. In all these primitives, secret keys are represented as quantum states that retain the functionality of the original secret keys. We design revocation procedures and guarantee that once a user successfully passes the procedure, they cannot compute the functionality any more.

All our constructions are secure under the quantum subexponential hardness of learning with errors [Reg05]—provided that revocation succeeds with high probability. At the heart of all of our contributions lies our result which shows that the Dual-Regev public-key encryption scheme of [GPV07] satisfies revocation security.

Related Notions. There are several recent notions in quantum cryptography that are related to revocability. Of particular relevance is the stronger notion of copy-protection introduced by Aaronson [Aar09]. Breaking the revocable security of a task gives the adversary a way to make two copies of a (possibly different) state both of which are capable of computing the same functionality. Thus, copy-protection is a stronger notion. However, the only known constructions of copy-protection schemes [CLLZ21, LLQZ22] rely on the heavy hammer of *post-quantum secure* indistinguishability obfuscation for which there are no known constructions based on well-studied assumptions. Our constructions, in contrast, rely on the post-quantum hardness of the standard learning with errors problem. A different related notion is the significantly weaker definition of secure software leasing [AL21] which guarantees that once the quantum state computing a functionality is returned, the *honest evaluation algorithm* cannot compute the original functionality. Yet another orthogonal notion is that of certifiably deleting *ciphertexts*, originating from the works of Unruh [Unr13] and Broadbent and Islam [BI20b]. In contrast, our goal is to delegate and revoke *cryptographic capabilities* enabled by private keys. For detailed comparisons, we refer the reader to Section 1.5.

1.1 Our Contributions in More Detail

We present our results in more detail below. First, we introduce the notion of key-revocable public-key encryption. Our main result is that the Dual-Regev public-key encryption scheme [GPV07] satisfies revocation security. After that, we study revocation security in the context of fully homomorphic encryption and pseudorandom functions.

Key-Revocable Public-Key Encryption. We consider public-key encryption schemes where the decryption key, modeled as a quantum state, can be delegated to a third party and can later be revoked [GZ20]. The syntax of a key-revocable public-key scheme (Definition 5.1) is as follows:

- $\mathsf{KeyGen}(1^\lambda)$: this is a setup procedure which outputs a public key PK, a master secret key MSK and a decryption key ρ_{SK}. While the master secret key is typically a classical string, the

[1] The attack would indeed have gone undetected but for the Snowden revelations.

decryption key is modeled as a quantum state. (The use cases of MSK and ρ_{SK} are different, as will be evident below.)

- $\mathsf{Enc}(\mathsf{PK}, x)$: this is the regular classical encryption algorithm which outputs a ciphertext CT.
- $\mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT})$: this is a quantum algorithm which takes as input the quantum decryption key ρ_{SK} and a classical ciphertext, and produces a plaintext.
- $\mathsf{Revoke}(\mathsf{PK}, \mathsf{MSK}, \sigma)$: this is the revocation procedure that outputs Valid or $\mathsf{Invalid}$. If σ equals the decryption key ρ_{SK}, then Revoke outputs Valid with high probability.

After the decryption key is returned, we require that the sender loses its ability to decrypt ciphertexts. This is formalized as follows (see Definition 5.3): conditioned on revocation being successful, no adversary can distinguish whether it is given an encryption of a message versus the uniform distribution over the ciphertext space with advantage better than $\mathsf{negl}(\lambda)$. Moreover, we require that revocation succeeds with a probability negligibly close to 1 (more on this later).
We prove the following in Theorem 6.2.

Theorem (Informal)**.** *Assuming that the* LWE *and* SIS *problems with subexponential modulus are hard against quantum adversaries running in subexponential time (see Section 2.2), there exists a key-revocable public-key encryption scheme.*

Due to the quantum reduction from SIS to LWE [SSTX09], the two assumptions are, in some sense, equivalent. Therefore, we can in principle rely on the subexponential hardness of LWE alone.

Our work improves upon prior works, which either use post-quantum secure indistinguishability obfuscation [GZ20, CLLZ21] or consider the weaker private-key setting [KN22].

Key-Revocable Public-Key Encryption with Classical Revocation. Our previous notion of key-revocable encryption schemes models the key as a quantum state which can later be "returned." One may therefore reasonably ask whether it is possible to achieve a classical notion of revocation, similar to the idea of deletion certificates in the context of quantum encryption [BI20b, HMNY21b, BK22, Por22, HMNY21a]. Rather then return a quantum decryption key, the recipient would simply apply an appropriate measurement (as specified by a procedure Delete), and output a classical certificate which can be verified using Revoke. We also consider key-revocable public-key encryption scheme with classical revocation (see Definition 7.1) which has the same syntax as a regular key-revocable public-key scheme, except that revocation process occurs via the following two procedures:

- $\mathsf{Delete}(\rho_{\mathsf{SK}})$: this takes as input a quantum decryption key ρ_{SK}, and produces a classical revocation certificate π.
- $\mathsf{Revoke}(\mathsf{PK}, \mathsf{MSK}, \pi)$: this takes as input the (classical) master secret key MSK and a (classical) certificate π, and outputs Valid or $\mathsf{Invalid}$. If π is the output of $\mathsf{Delete}(\rho_{\mathsf{SK}})$, then Revoke outputs Valid with high probability.

Our notion of security (see Definition 7.2) is essentially the same as for a regular key-revocable public-key encryption scheme where revocation is quantum. We prove the following in Theorem 7.13.

Theorem (Informal)**.** *Assuming that the* LWE *and* SIS *problems with subexponential modulus are hard against quantum adversaries running in subexponential time (see Section 2.2), there exists a key-revocable public-key encryption scheme with classical revocation.*

The assumptions required for the theorem are essentially the same as for our previous Dual-Regev scheme.

Key-Revocable Fully Homomorphic Encryption. We go beyond the traditional public-key setting and design the first *fully homomorphic encryption* (FHE) scheme [Gen09, BV14b] with key-revocation capabilities. Our construction is based on a variant of the (leveled) FHE scheme of Gentry, Sahai and Waters [GSW13], which we extend to a key-revocable encryption scheme using Gaussian superpositions. The syntax of a key-revocable FHE scheme is the same as in the key-revocable public-key setting from before (Definition 5.1), except for the additional algorithm Eval which is the same as in a regular FHE scheme. We prove the following in Theorem 8.3.

Theorem (Informal)**.** *Assuming that the* LWE *and* SIS *problems with subexponential modulus are hard against quantum adversaries running in subexponential time (see Section 2.2), there exists a key-revocable (leveled) fully homomorphic encryption scheme.*

We prove the theorem by invoking the security of our key-revocable Dual-Regev public-key encryption scheme in Section 6. Similar to the case of revocable PKE with classical revocation, we can also consider revocable FHE with classical revocation, which can be achieved based on the same assumptions as above.

(Key-)Revocable Pseudorandom Functions. We consider other cryptographic primitives with key-revocation capabilities that go beyond decryption functionalities; specifically, we introduce the notion of *key-revocable* pseudorandom functions (PRFs) with the following syntax:

- $\mathsf{Gen}(1^\lambda)$: outputs a PRF key k, a quantum key ρ_k and a master secret key MSK.
- $\mathsf{PRF}(k; x)$: on key k and input x, output a value y. This is a deterministic algorithm.
- $\mathsf{Eval}(\rho_k, x)$: on input a state ρ_k and an input x, output a value y.
- $\mathsf{Revoke}(\mathsf{MSK}, \sigma)$: on input a verification key MSK and state σ, outputs Valid or Invalid.

After the quantum key ρ_k is successfully returned, we require that the sender loses its ability to evaluate the PRF. This is formalized as follows (see Definition 9.3): no efficient adversary can simultaneously pass the revocation phase and succeed in distinguishing the output of a pseudorandom function on a random challenge input x^* versus uniform with advantage better than $\mathsf{negl}(\lambda)$. In fact, we consider a more general definition where the adversary receives many challenge inputs instead of just one.

We give the first construction of key-revocable pseudorandom functions (PRFs) from standard assumptions. Previous schemes implicit in [CLLZ21] either require indistinguishability obfuscation, or considered weaker notions of revocable PRFs in the form of *secure software leasing* [AL21, KNY21a], which merely prevents the possiblity of *honestly* evaluating the PRF once the key is revoked.

Since in the context of pseudorandom functions, it is clear what is being revoked, we instead simply call the notion revocable pseudorandom functions.

Theorem (Informal)**.** *Assuming that the* LWE *and* SIS *problems with subexponential modulus are hard against quantum adversaries running in subexponential time (see Section 2.2), there exist key-revocable pseudorandom functions.*

Revocable pseudorandom functions immediately give us key-revocable (many-time secure) secret-key encryption schemes. We also revocable PRFs with classical revocation can be achieved based on the same assumptions as above.

Inverse-polynomial revocation based on conjectures. In all the results above, we assume that the probability of revocation is negligibly close to 1. Even in this restrictive setting, our proofs turn out to be highly non-trivial and require careful use of a diverse set of techniques! Moreover, to date, no constructions of key-revocable PRFs or FHE were known based on assumptions weaker than post-quantum iO.

A natural question to explore is whether we can achieve the following stronger security notion of revocable public-key encryption: if the adversary successfully revokes the decryption key with inverse-polynomial probability, then semantic security of revocable PKE still holds. If we achieved this stronger notion of revocable PKE, then we would also achieve the corresponding stronger notions of revocable PRFs and FHE based on the same computational assumptions.

We show how to achieve these stronger results based on a plausible conjecture (see Conjecture 1) which we can phrase in the language of *asymmetric* cloning games [AKL23]. Informally, our conjecture states the following: suppose that our Dual-Regev PKE scheme is not key-revocable (according to the above stronger definition). Then, there exists a triplet of algorithms $(\mathcal{A}, \mathcal{B}, \mathcal{C})$, where

- $\mathcal{A}$ is a QPT algorithm which receives as input $(\mathbf{A}, \mathbf{y}, |\psi_{\mathbf{y}}\rangle)$, where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and $|\psi_{\mathbf{y}}\rangle$ is a Gaussian superposition of all the short vectors mapping $\mathbf{A}$ to $\mathbf{y}$. It produces a bipartite state on two registers B and C.
- $\mathcal{B}$ is a fixed and inefficient revocation algorithm which receives as input $(\mathbf{A}, \mathbf{y}, \mathrm{B})$ and applies the (inefficient) projective measurement $\left\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, I - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\right\}$ to register B.
- $\mathcal{C}$ is a QPT algorithm which on input $(\mathbf{A}, \mathbf{y}, \mathrm{C})$ distinguishes Dual-Regev ciphertexts (with respect to $\mathbf{A}$ and $\mathbf{y}$) from random[2] with inverse-polynomial advantage—conditioned on the event that the revocation algorithm $\mathcal{B}$ succeeds on register B.

The difficulty in proving the conjecture lies in the fact that one needs to invoke the LWE assumption with respect to $\mathcal{C}$ who holds C, while at the same time guaranteeing that an inefficient projective measurement succeeds on a separate register B. We leave proving (or refuting) the above conjecture to future works.

We also consider another variant of the above conjecture (see Conjecture 2) in the context of classical revocation. Here, $\mathcal{A}$ is instead given $|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle$, which is essentially $|\psi_{\mathbf{y}}\rangle$, except that a phase of $\omega_q^{\langle \mathbf{x}, \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}\rangle}$ is planted for the term $|\mathbf{x}\rangle$ in the superposition. Moreover, we replace $\mathcal{B}$ with a fixed classical revocation algorithm which can detect $\mathbf{v}$ within ciphertexts $\mathsf{CT} \approx \mathbf{s}^\intercal\mathbf{A} + \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}^\intercal$, for some $\mathbf{s} \in \mathbb{Z}_q^n$, given the register B and an appropriate *trapdoor* for $\mathbf{A}$.

[2]Strictly speaking, the challenge input here consists of a vector which resembles a Dual-Regev ciphertext but which implicitly also depends on a fixed Gaussian vector $\mathbf{x}_0$ such that $\mathbf{A}\mathbf{x}_0 = \mathbf{y} \pmod q$. More on this in Section 1.2.

Discussion: Unclonable Cryptography from LWE. Over the years, the existence of many fundamental cryptographic primitives such as pseudorandom functions [BPR12], fully homomorphic encryption [BV14a], attribute-based encryption [BGG+14] and succinct argument systems [CJJ22] have been based on the existence of learning with errors. In fact, as far as we know, there are only a few foundational primitives remaining (indistinguishability obfuscation is one such example) whose existence is not (yet) known to be based on learning with errors.

The situation is quite different in the world of unclonable cryptography. Most of the prominent results have information-theoretic guarantees but restricted functionalities [BI20b, BL20] or are based on the existence of post-quantum indistinguishability obfuscation [Zha21, CLLZ21]. While there are works [KNY21b] that do propose lattice-based constructions of unclonable primitives, there are still many primitives, such as quantum money and quantum copy-protection, whose feasibility we would like to establish based on the post-quantum hardness of the learning with errors (or other such relatively well-studied assumptions). We hope that our work presents a new toolkit towards building more unclonable primitives from LWE.

Independent and Concurrent Work. Independently and concurrently[3], Agrawal et al. [AKN+23] explored the notion of public-key encryption with secure leasing which is related to key-revocable public-key encryption. We highlight the differences below:

- *Advanced notions*: We obtain key-revocable *fully homomorphic encryption* and key-revocable *pseudorandom functions* which are unique to our work. They explore other notions of advanced encryption with secure leasing including attribute-based encryption and functional encryption, which are not explored in our work.
- *Security definition*: We consider the stronger notion of classical revocation[4] whereas they consider quantum revocation.
- *Public-key encryption*: They achieve a generic construction based on any post-quantum secure public-key encryption[5] whereas our notion is based on the post-quantum hardness of the learning with errors problem, or on newly introduced conjectures (when revocation does not succeed with high probability). Their construction of revocable public-key encryption involves many complex abstractions whereas our construction is based on the versatile Dual-Regev public-key encryption scheme.

1.2 Overview

We now give a technical overview of our constructions and their high level proof ideas. We begin with the key-revocable public-key encryption construction. A natural idea would be to start with Regev's public-key encryption scheme [Reg05] and to then upgrade the construction in order to make it revocable. However, natural attempts to associate an unclonable quantum state with the decryption key fail and thus, we instead consider the Dual-Regev public-key encryption scheme and make it key-revocable. We describe the scheme below.

[3] Both of our works were posted online around the same time.

[4] The stronger notion was updated in our paper subsequent to posting of our and their work.

[5] Their construction achieves the stronger definition where the revocation only needs to succeed with inverse polynomial probability.

Key-Revocable Dual-Regev Public-Key Encryption. Our first construction is based on the *Dual-Regev* public-key encryption scheme [GPV07] and makes use of Gaussian superpositions which serve as a quantum decryption key. We give an overview of Construction 1 below.

- $\mathsf{KeyGen}(1^n)$: sample a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ along with a *short trapdoor basis* $\mathsf{td}_{\mathbf{A}}$. To generate the decryption key, we employ the following procedure[6]: Using the matrix $\mathbf{A}$ as input, first create a Gaussian superposition of short vectors in $\mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$, denoted by[7]

$$|\psi\rangle = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \otimes |\mathbf{A}\cdot\mathbf{x} \pmod q\rangle$$

  where $\rho_\sigma(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/\sigma^2)$ is the Gaussian measure, for some $\sigma > 0$. Next, measure the second register which partially collapses the superposition and results in the *coset state*

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle$$

  for some outcome $\mathbf{y} \in \mathbb{Z}_q^n$. Finally, we let $|\psi_{\mathbf{y}}\rangle$ be the decryption key ρ_{SK}, $(\mathbf{A}, \mathbf{y})$ be the public key PK, and we let the trapdoor $\mathsf{td}_{\mathbf{A}}$ serve as the master secret key MSK.

- $\mathsf{Enc}(\mathsf{PK}, \mu)$ is identical to Dual-Regev encryption. To encrypt a bit $\mu \in \{0,1\}$, sample a random string $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ together with discrete Gaussian errors $\mathbf{e} \in \mathbb{Z}^m$ and $e' \in \mathbb{Z}$, and output the (classical) ciphertext CT given by

$$\mathsf{CT} = (\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal, \mathbf{s}^\intercal\mathbf{y} + e' + \mu\cdot\lfloor\frac{q}{2}\rfloor) \quad \in \mathbb{Z}_q^m \times \mathbb{Z}_q.$$

- $\mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT})$: to decrypt a ciphertext CT using the decryption key $\rho_{\mathsf{SK}} = |\psi_{\mathbf{y}}\rangle$, first apply the unitary $U : |\mathbf{x}\rangle\,|0\rangle \to |\mathbf{x}\rangle\,|\mathsf{CT}\cdot(-\mathbf{x}, 1)^\intercal\rangle$ on input $|\psi_{\mathbf{y}}\rangle\,|0\rangle$, and then measure the second register in the computational basis. Because $|\psi_{\mathbf{y}}\rangle$ is a superposition of short vectors $\mathbf{x}$ subject to $\mathbf{A}\cdot\mathbf{x} = \mathbf{y} \pmod q$, we obtain an approximation of $\mu\cdot\lfloor\frac{q}{2}\rfloor$ from which we can recover μ.[8]

- $\mathsf{Revoke}(\mathsf{PK}, \mathsf{MSK}, \rho)$: to verify the returned state ρ given as input the public key $(\mathbf{A}, \mathbf{y})$ and master secret key $\mathsf{td}_{\mathbf{A}}$, apply the projective measurement $\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, I - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\}$ onto ρ. Output Valid, if the measurement succeeds, and output $\mathsf{Invalid}$, otherwise.

Implementing revocation, efficiently. Note that performing a projective measurement onto a fixed Gaussian state $|\psi_{\mathbf{y}}\rangle$ is, in general, computationally infeasible. In fact, if it were to be possible to efficiently perform this projection using $(\mathbf{A}, \mathbf{y})$ alone, then one could easily use such a procedure to solve the short integer solution (SIS) problem. Fortunately, we additionally have the trapdoor for $\mathbf{A}$ at our disposal in order to perform such a projection.

One of our contributions is to design a *quantum discrete Gaussian sampler for q-ary lattices*[9] which, given as input $(\mathbf{A}, \mathbf{y}, \mathsf{td}_{\mathbf{A}}, \sigma)$, implements a unitary that efficiently prepares the Gaussian

[6] In Section 3.2, this is formalized as the procedure $\mathsf{GenGauss}$ (see Algorithm 1).

[7] Note that the state is not normalized for convenience.

[8] For approriate choices of parameters, decryption via rounding succeeds at outputting μ with overwhelming probability and hence we can invoke the "almost as good as new" lemma [Aar16] to recover the original state $|\psi_{\mathbf{y}}\rangle$.

[9] In Section 3.3, this is formalized as the procedure $\mathsf{QSampGauss}$ (see Algorithm 2).

superposition $|\psi_{\mathbf{y}}\rangle$ from scratch with access to the trapdoor $\mathsf{td}_{\mathbf{A}}$. At a high level, our Gaussian sampler can be alternately thought of as an explicit quantum reduction from the *inhomogenous* SIS problem [Ajt96] to the search variant of the LWE problem (see Section 3.3).

Insight: Reduction to SIS. Our goal is to use the state returned by the adversary and to leverage the indistinguishability guarantee in order to break some computational problem. It should seem suspicious whether such a reduction is even possible: after all the adversary is returning the state we gave them! *How could this possibly help?* Our main insight lies in the following observation: while the adversary does eventually return the state we give them, the only way it can later succeed in breaking the semantic security of dual Regev PKE is if it retains useful information about the state. If we could somehow extract this information from the adversary, then using the extracted information alongside the returned state, we could hope to break some computational assumption. For instance, suppose we can extract a short vector $\mathbf{x}$ such that $\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod q$. By measuring the state returned by the adversary, we could then hope to get a second short vector $\mathbf{x}'$ such that $\mathbf{A} \cdot \mathbf{x}' = \mathbf{y} \pmod q$, and from this, we can recover a short solution in the kernel of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$.

Even if, for a moment, we disregard the issue of being able to extract $\mathbf{x}$ from the adversary, there are still some important missing steps in the above proof template:

- Firstly, measuring the returned state should give a vector different from $\mathbf{x}$ with non-negligible probability. In order to prove this, we need to argue that the squared ampltidue of every term is bounded away from 1. We prove this statement (Lemma 2.10) holds as long as $\mathbf{A}$ is full rank.

- Secondly, the reduction to SIS would only get as input $\mathbf{A}$ and not a trapdoor for $\mathbf{A}$. This means that it will no longer be possible for the reduction to actually check whether the state returned by the adversary is valid. We observe that, instead of first verifying whether the returned state is valid and then measuring in the computational basis, we can in fact skip verification and immediately go ahead and measure the state in the computational basis; this is implicit in the analysis in the proof of Lemma 6.20.

- Finally, the adversary could have entangled the returned state with its residual state in such a way that measuring the returned state always yields the same vector $\mathbf{x}$ as the one extracted from the adversary. In the same analysis in the proof of Lemma 6.20, we prove that, even if the adversary entangles its state with the returned state, with non-negligible probability we get two distinct short vectors mapping $\mathbf{A}$ to $\mathbf{y}$.

All that is left is to argue that it is possible to extract $\mathbf{x}$ from the adversary while simultaneously verifying whether the returned state is correct or not. To show that we can indeed extract another short pre-image from the adversary's quantum side information, we make use of what we call a *simultaneous search-to-decision reduction with quantum auxiliary input* for the Dual-Regev scheme.

Main contribution: Simultaneous search-to-decision reduction with quantum advice. Informally, our theorem says the following: any successful Dual-Regev distinguisher with access to quantum side information Aux (which depends on the decryption key) can be converted into a successful extractor that finds a key on input Aux – even conditioned on Revoke succeeding on a seperate register R. We now present some intuition behind our proof.

Suppose there exists a successful Dual-Regev distinguisher $\mathcal{D}$ (as part of the adversary $\mathcal{A}$) that, given quantum auxiliary information Aux, can distinguish between $(\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal, \mathbf{s}^\intercal\mathbf{y}+e')$ and uniform $(\mathbf{u}, r) \in \mathbb{Z}_q^m \times \mathbb{Z}_q$ with advantage ϵ.

Ignoring register R: For now, let us ignore the fact that Revoke is simultaneously applied on system R. Inspired by techniques from the *leakage resilience* literature [DGT+10], we now make the following observation. Letting $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}_0 \pmod q$, for some Gaussian vector $\mathbf{x}_0$ with distribution proportional to $\rho_\sigma(\mathbf{x}_0)$, the former sample can be written as $(\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal, (\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal) \cdot \mathbf{x}_0 + e')$. Here, we assume a *noise flooding* regime in which the noise magnitude of e' is significantly larger than that of $\mathbf{e}^\intercal \cdot \mathbf{x}_0$. Because the distributions are statistically close, the distinguisher $\mathcal{D}$ must succeed at distinguishing the sample from uniform with probability negligibly close to ϵ. Finally, we invoke the LWE assumption and claim that the same distinguishing advantage persists, even if we replace $(\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal)$ with a random string $\mathbf{u} \in \mathbb{Z}_q^m$. Here, we rely on the fact that the underlying LWE sample is, in some sense, independent of the auxiliary input Aux handed to the distinguisher $\mathcal{D}$. To show that this is the case, we need to argue that the reduction can generate the appropriate inputs to $\mathcal{D}$ on input $\mathbf{A}$; in particular it should be able to generate the auxiliary input Aux (which depends on a state $|\psi_\mathbf{y}\rangle$), while simultaneously producing a Gaussian vector $\mathbf{x}_0$ such that $\mathbf{A} \cdot \mathbf{x}_0 = \mathbf{y} \pmod q$. Note that this seems to violate the SIS assumption, since the ability to produce both a superposition $|\psi_\mathbf{y}\rangle$ of pre-images and a single pre-image $\mathbf{x}_0$ would allow one to obtain a collision for $\mathbf{y}$.

Invoking Gaussian-collapsing: To overcome this issue, we ask the reduction to generate the quantum auxiliary input in a different way; rather than computing Aux as a function of $|\psi_\mathbf{y}\rangle$, we compute it as a function of $|\mathbf{x}_0\rangle$, where $\mathbf{x}_0$ results from *collapsing* the state $|\psi_\mathbf{y}\rangle$ via a measurement in the computational basis. By invoking the *Gaussian collapsing property* [Por22], we can show that the auxiliary information computed using $|\psi_\mathbf{y}\rangle$ is computationally indistinguishable from the auxiliary information computed using $|\mathbf{x}_0\rangle$. Once we invoke the collapsed version of $|\psi_\mathbf{y}\rangle$, we can carry out the reduction and conclude that $\mathcal{D}$ can distinguish between the samples $(\mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0)$ and $(\mathbf{u}, r)$, where $\mathbf{u}$ and r are random and $\mathbf{x}_0$ is Gaussian, with advantage negligibly close to ϵ.[10] Notice that $\mathcal{D}$ now resembles a so-called *Goldreich-Levin* distinguisher [GL89].

Reduction to Goldreich-Levin: Assuming the existence of a quantum Goldreich-Levin theorem for the field $\mathbb{Z}_q$, one could then convert $\mathcal{D}$ into an extractor that extracts $\mathbf{x}_0$ with high probability. Prior to our work, a quantum Goldreich-Levin theorem was only known for $\mathbb{Z}_2$ [AC02, CLLZ21]. In particular, it is unclear how to extend prior work towards higher order fields $\mathbb{Z}_q$ because the interference pattern in the analysis of the quantum extractor does not seem to generalize beyond the case when $q = 2$. Fortunately, we can rely on the *classical* Goldreich-Levin theorem for finite fields due to Dodis et al. [DGT+10], as well as recent work by Bitansky, Brakerski and Kalai. [BBK22] which shows that a large class of classical reductions can be generically converted into a quantum reductions. This allows us to obtain a quantum Goldreich-Levin theorem for large fields, which we prove in Section 4. Specifically, we can show that a distinguisher $\mathcal{D}$ that, given auxiliary input Aux, can distinguish between $(\mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0)$ and $(\mathbf{u}, r)$ with advantage ε can be converted into a quantum extractor that can extract $\mathbf{x}_0$ given Aux in time $\mathrm{poly}(1/\varepsilon, q)$ with probability $\mathrm{poly}(\varepsilon, 1/q)$.

[10] Technically, $\mathcal{D}$ can distinguish between $(\mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 + e')$ and $(\mathbf{u}, r)$ for a Gaussian error e'. However, by defining a distinguisher $\tilde{\mathcal{D}}$ that first shifts $\mathbf{u}$ by a Gaussian vector e' and then runs $\mathcal{D}$, we obtain the desired distinguisher.

Incorporating the revoked register R: To complete the security proof behind our key-revocable Dual-Regev scheme, we need to show something *stronger*; namely, we need to argue that the Goldreich-Levin extractor succeeds on input Aux – even conditioned on the fact that Revoke outputs Valid when applied on a separate register R (which may be entangled with Aux). We can consider two cases based on the security definition.

- Revocation succeeds with probability negligibly close to 1: in this case, applying the revocation or not does not make a difference since the state before applying revocation is negligibly close (in trace distance) to the state after applying revocation. Thus, the analysis is essentially the same as the setting where we ignore the register R.
- Revocation is only required to succeed with probability $1/\mathrm{poly}(\lambda)$: in this case, we do not know how to formally prove that the extractor and Revoke simultaneously succeed with probability $1/\mathrm{poly}(\lambda)$. Thus, we state this as a conjecture (see Theorem 6.16) and leave the investigation of this conjecture to future works.

1.3 Key-Revocable Dual-Regev Encryption with Classical Revocation

Recall that our key-revocable Dual-Regev public-key encryption scheme requires that a quantum state is *returned* as part of revocation. In Construction 2, we give a Dual-Regev encryption scheme with *classical key revocation*. The idea behind our scheme is the following: We use the same Gaussian decryption key from before, except that we also plant an appropriate phase into the state

$$|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle = \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y}}} \rho_\sigma(\mathbf{x})\, \omega_q^{\langle \mathbf{x}, \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}\rangle} |\mathbf{x}\rangle\,,$$

where $\mathbf{Z}_q$ is the generalized q-ary Pauli-Z operator, $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$ is a random string and ν is a sufficiently large integer with $\nu \ll q$ (to be determined later). The reason the complex phase is useful is the following: To revoke, we simply ask the recipient of the state to apply the Fourier transform and to measure in the computational basis. This results in a shifted LWE sample

$$\mathbf{w} = \hat{\mathbf{s}}^\intercal \mathbf{A} + \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}^\intercal + \hat{\mathbf{e}}^\intercal \in \mathbb{Z}_q^m,$$

where $\hat{\mathbf{s}}$ is random and $\hat{\mathbf{e}}$ is Gaussian. Note that $\mathbf{w}$ can easily be decrypted with access to a short trapdoor basis for $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$. In particular, we will accept the *classical* revocation certificate $\mathbf{w}$ if and only if it yields the string $\mathbf{v}$ once we decrypt it. This way there is no need to return a quantum state as part of revocation, and we can essentially "force" the adversary to forget the decryption key. Here, we crucially rely on the fact that Dual-Regev decryption keys are encoded in the *computational basis*, whereas the revocation certificate can only be obtained via a measurement in the incompatible *Fourier basis*. Intuitively, it seems impossible for any computationally bounded adversary to recover both pieces of information *at the same time*.

To prove security, we follow a similar proof as in our previous construction from Section 6. First, we observe that any successful distinguisher that can distinguish Dual-Regev from uniform with quantum auxiliary input Aux (once revocation has taken place) can be converted into an extractor that obtains a valid pre-image of $\mathbf{y}$. Here, we use our Goldreich-Levin search-to-decision reduction with quantum auxiliary input. However, because the information returned to the challenger is

not in fact a superposition of pre-images anymore, we cannot hope to break SIS as we did before. Instead, we make the following observation: if the adversary simultaneously succeeds at passing revocation as well as distinguishing Dual-Regev ciphertexts from uniform, this means the adversary

- knows an LWE encryption $\mathbf{w} \in \mathbb{Z}_q^m$ that yields $\mathbf{v} \in \{0,1\}^m$ when decrypted using a short trapdoor basis for $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and
- can extract a short pre-image of $\mathbf{y}$ from its auxiliary input Aux.

In other words, the adversary simultaneously knows information about the computational basis as well as the Fourier domain of the Gaussian superposition state. We show that this violates a generic collapsing-type property of the Ajtai hash function recently proven by Bartusek, Khurana and Poremba [BKP23, Theorem 5.5], which relies on the hardness of subexponential LWE and SIS.

1.4 Applications

We leverage our result of key-revocable Dual-Regev encryption to get key-revocable fully homomorphic encryption and revocable pseudorandom functions. While our constructions can easily be adapted to enable *classical revocation* (via our Dual-Regev scheme with classical key revocation in Construction 2), we choose to focus on the quantum revocation setting for simplicity.

Key-Revocable Dual-Regev Fully Homomorphic Encryption. Our first application of our key-revocable public-key encryption concerns fully homomorphic encryption schemes. We extend our key-revocable Dual-Regev scheme towards a (leveled) FHE scheme in Construction 3 by using the DualGSW variant of the FHE scheme by Gentry, Sahai and Waters [GSW13, Mah18].

To encrypt a bit $\mu \in \{0,1\}$ with respect to the public-key $(\mathbf{A}, \mathbf{y})$, sample a matrix $\mathbf{S} \xleftarrow{\$} \mathbb{Z}_q^{n \times N}$ together with a Gaussian error matrix $\mathbf{E} \in \mathbb{Z}^{m \times N}$ and row vector $\mathbf{e} \in \mathbb{Z}^N$, and output the ciphertext

$$\mathsf{CT} = \left[\frac{\mathbf{A}^\intercal\mathbf{S}+\mathbf{E}}{\mathbf{y}^\intercal\mathbf{S}+\mathbf{e}}\right] + \mu \cdot \mathbf{G} \pmod q \in \mathbb{Z}_q^{(m+1)\times N}.$$

Here, $\mathbf{G}$ is the *gadget matrix* which converts a binary vector in into its field representation over $\mathbb{Z}_q$. As before, the decryption key consists of a Gaussian superposition $|\psi_{\mathbf{y}}\rangle$ of pre-images of $\mathbf{y}$.

Note that the DualGSW ciphertext can be thought of as a column-wise concatenation of N-many independent Dual-Regev ciphertexts. In Theorem 8.3, we prove the security of our construction by invoking the security of our key-revocable Dual-Regev scheme.

Revocable Pseudorandom Functions. Our next focus is on leveraging the techniques behind key-revocable public-key encryption to obtain revocable pseudorandom functions. Recall that the revocation security of pseudorandom functions stipulates the following: any efficient adversary (after successfully revoking the state that enables it to evaluate pseudorandom functions) cannot distinguish whether it receives pseudorandom outputs on many challenge inputs versus strings picked uniformly at random with advantage better than $\mathsf{negl}(\lambda)$. An astute reader might notice that revocation security does not even imply the traditional pseudorandomness guarantee! Hence, we need to additionally impose the requirement that a revocable pseudorandom function should also satisfy the traditional pseudorandomness guarantee.

Towards realizing a construction satisfying our definitions, we consider the following template:

1. First show that there exists a μ-revocable pseudorandom function for $\mu = 1$. Here, μ-revocation security means the adversary receives μ-many random inputs after revocation.

2. Next, we show that any 1-revocable pseudorandom function also satisfies the stronger notion of revocation security where there is no a priori bound on the number of challenge inputs received by the adversary.

3. Finally, we show that we can generically upgrade any revocable PRF in such a way that it also satisfies the traditional pseudorandomness property.

The second bullet is proven using a hybrid argument. The third bullet is realized by combining a revocable PRF with a post-quantum secure PRF (not necessarily satisfying revocation security).

Hence, we focus the rest of our attention on proving the first bullet.

1-revocation security. We start with the following warmup construction. The secret key k comprises of matrices $\mathbf{A}, \{\mathbf{S}_{i,0}, \mathbf{S}_{i,1}\}_{i\in[\ell],b\in\{0,1\}}$, where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, $\mathbf{S}_{i,b} \in \mathbb{Z}_q^{n\times n}$ such that all $\mathbf{S}_{i,b}$ are sampled from some error distribution and the output of the pseudorandom function on x is denoted to be $\lfloor \sum_{i\in[\ell]} \mathbf{S}_{i,x_i}\mathbf{A}\rceil_p$, where $q \gg p$ and $\lfloor\cdot\rceil_p$ refers to a particular rounding operation modulo p.

In addition to handing out a regular PRF key k, we also need to generate a quantum key ρ_k such that, given ρ_k and any input x, we can efficiently compute $\mathsf{PRF}(k, x)$. Moreover, ρ_k can be revoked such that any efficient adversary after revocation loses the ability to evaluate the pseudorandom function. To enable the generation of ρ_k, we first modify the above construction. We generate $\mathbf{y} \in \mathbb{Z}_q^n$ and include this as part of the key. The modified pseudorandom function, on input x, outputs $\lfloor \sum_{i\in[\ell]} \mathbf{S}_{i,x_i}\mathbf{y}\rceil_p$. We denote $\sum_{i\in[\ell]} \mathbf{S}_{i,x_i}$ by $\mathbf{S}_x$ and, with this new notation, the output of the pseudorandom function can be written as $\lfloor \mathbf{S}_x\mathbf{y}\rceil_p$.

With this modified construction, we now describe the elements as part of the quantum key ρ_k:

- For every $i \in [\ell]$, include $\mathbf{S}_{i,b}\mathbf{A} + \mathbf{E}_{i,b}$ in ρ_k, where $i \in [\ell]$ and $b \in \{0, 1\}$. We sample $\mathbf{S}_{i,b}$ and $\mathbf{E}_{i,b}$ from a discrete Gaussian distribution with appropriate standard deviation $\sigma > 0$.

- Include $|\psi_{\mathbf{y}}\rangle$ which, as defined in the key-revocable Dual-Regev construction, is a Gaussian superposition of short solutions mapping $\mathbf{A}$ to $\mathbf{y}$.

To evaluate on an input x using ρ_k, compute $\sum_i \mathbf{S}_{i,x_i}\mathbf{A} + \mathbf{E}_{i,x_i}$ and then using the state $|\psi_{\mathbf{y}}\rangle$, map this to $\sum_i \mathbf{S}_{i,x_i}\mathbf{y} + \mathbf{E}_{i,x_i}$. Finally, perform the rounding operation to get the desired result.

Our goal is to show that after the adversary revokes $|\psi_{\mathbf{y}}\rangle$, on input a challenge input x^* picked uniformly at random, it cannot predict whether it has received $\lfloor \sum_{i\in[N]} \mathbf{S}_{i,x_i^*}\mathbf{y}\rceil_p$ or a uniformly random vector in $\mathbb{Z}_p^n$.

Challenges in proving security: We would like to argue that when the state $|\psi_{\mathbf{y}}\rangle$ is revoked, the adversary loses its ability to evaluate the pseudorandom function. Unfortunately, this is not completely true. For all we know, the adversary could have computed the pseudorandom function on many inputs of its choice before the revocation phase and it could leverage this to break the security after revocation. For instance, suppose say the input is of length $O(\log\lambda)$ then in this case, the adversary could evaluate the pseudorandom function on all possible inputs before revocation. After revocation, on any challenge input x^*, the adversary can then successfully predict whether it receives a pseudorandom output or a uniformly chosen random output. Indeed, a pseudorandom function with $O(\log\lambda)$-length input is learnable and hence, there should be no hope of proving it

to be key-revocable. This suggests that, at the very least, we need to explicitly incorporate the fact that x^* is of length $\omega(\log \lambda)$, and more importantly, should have enough entropy, in order to prove security.

Our insight: Our insight is to reduce the security of revocable pseudorandom function to the security of key-revocable Dual-Regev public-key encryption. At a high level, our goal is to set up the parameters in such a way that the following holds:

- $(\mathbf{A}, \mathbf{y})$, defined above, is set to be the public key corresponding to the Dual-Regev public-key encryption scheme,
- $|\psi_{\mathbf{y}}\rangle$, which is part of the pseudorandom function key, is set to be the decryption key of the Dual Regev scheme,
- Suppose that the challenge ciphertext, denoted by CT^*, comprises of two parts: $\mathsf{CT}_1^* \in \mathbb{Z}_q^{n \times m}$ and $\mathsf{CT}_2^* \in \mathbb{Z}_q^n$. Note that if $\mathsf{CT}_1^* \approx \mathbf{s}^\intercal \mathbf{A}$ and $\mathsf{CT}_2^* \approx \mathbf{s}^\intercal \mathbf{y}$, for some LWE secret vector $\mathbf{s}$, then CT_1^* can be mapped onto CT_2^* using the state $|\psi_{\mathbf{y}}\rangle$. We use CT_1^* to set the challenge input x^* in such a way that CT_2^* is the output of the pseudorandom function on x^*. This implicitly resolves the entropy issue we discussed above; by the semantic security of Dual-Regev, there should be enough entropy in CT_1^* which translates to the entropy of x^*.

It turns out that our goal is quite ambitious: in particular, it is unclear how to set up the parameters such that the output of the pseudorandom function on x is exactly CT_2^*. Fortunately, this will not be a deterrant, we can set up the parameters such that the output is $\approx \mathsf{CT}_2^* + \mathbf{u}$, where $\mathbf{u}$ is a vector that is known to reduction.

Once we set up the parameters, we can then reduce the security of revocable pseudorandom functions to revocable Dual Regev.

Implementation details: So far we established the proof template should work but the implementation details of the proof need to be fleshed out. Firstly, we set up the parameters in such a way that $\ell = nm\lceil \log q \rceil$. This means that there is a bijective function mapping $[n] \times [m] \times [\lceil \log q \rceil]$ to $[\ell]$. As a result, the quantum key ρ_k can be alternately viewed as follows:

- For every $i \in [n], j \in [m], \tau \in [\lceil \log q \rceil], b \in \{0,1\}$, include $\mathbf{S}_b^{(i,j,\tau)} \mathbf{A} + \mathbf{E}_b^{(i,j,\tau)}$ in ρ_k. We sample $\mathbf{S}_b^{(i,j,\tau)}$ and $\mathbf{E}_b^{(i,j,\tau)}$ from a discrete Gaussian with appropriate standard deviation $\sigma > 0$.

The output of the pseudorandom function on input x can now be interpreted as

$$\mathsf{PRF}(k, x) = \left\lfloor \sum_{\substack{i \in [n], j \in [m] \\ \tau \in [\lceil \log q \rceil]}} \mathbf{S}_{x_i}^{(i,j,\tau)} \mathbf{y} \right\rceil_p$$

Next, we modify ρ_k as follows: instead of generating, $\mathbf{S}_b^{(i,j,\tau)} \mathbf{A} + \mathbf{E}_b^{(i,j,\tau)}$, we instead generate $\mathbf{S}_b^{(i,j,\tau)} \mathbf{A} + \mathbf{E}_b^{(i,j,\tau)} + \mathsf{M}_b^{(i,j,k)}$, for any set of matrices $\{\mathsf{M}_b^{(i,j,\tau)}\}$. The change should be undetectable to a computationally bounded adversary, thanks to the quantum hardness of learning with errors. In the security proof, we set up the challenge input x^* in such a way that summing up the matrices

$\mathsf{M}_{x_i^*}^{(i,j,\tau)}$ corresponds to CT_1^*, where CT_1^* is part of the key-revocable Dual-Regev challenge ciphertext as discussed above. With this modification, when ρ_k is evaluated on x^*, we get an output that is close to $\mathsf{CT}_2^* + \mathbf{u}$, where $\mathbf{u} \approx \sum_{i\in[n],j\in[m],\tau\in[\lceil\log(q)\rceil]} \mathbf{S}_{x_i}^{(i,j,\tau)}\mathbf{y}$ is known to the reduction (as discussed above) – thereby violating the security of key-revocable Dual-Regev scheme.

1.5 Related Work

Copy-Protection. Of particular relevance to our work is the notion of copy-protection 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.

While copy-protection is known to be impossible for arbitrary unlearnable functions [AL21, AK22], identifying interesting functionalities for which copy-protection is feasible has been an active research direction [CMP20, AKL+22, AKL23]. Of particular significance is the problem of copy-protecting cryptographic functionalities, such as decryption and signing functionalities. Coladangelo et al. [CLLZ21] took the first step in this direction and showed that it is feasible to copy-protect decryption functionalities and pseudorandom functions assuming the existence of post-quantum indistinguishability obfuscation. While a very significant first step, the assumption of post-quantum iO is unsatisfactory: there have been numerous post-quantum iO candidate proposals [BMSZ16, CVW18, BDGM20, DQV+21, GP21, WW21], but not one of them have been based on well-studied assumptions[11].

Our work can be viewed as copy-protecting cryptographic functionalities based on the post-quantum hardness of the learning with errors problem under a weaker yet meaningful security guarantee.

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[12], the honest evaluation algorithm on the residual state cannot compute the original functionality. Key-revocable encryption can be viewed as secure software leasing for decryption algorithms. However, unlike secure software leasing, key-revocable encryption satisfies a much stronger security guarantee, where there is no restriction on the adversary to run honestly after returning back the software. Secure leasing for different functionalities, namely, point functions [CMP20, BJL+21], evasive functions [AL21, KNY21b] and pseudorandom functions [ALL+21] have been studied by recent works.

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. Unruh's scheme uses conjugate coding [Wie83, BB84] and relies on the *monogamy of entanglement* in order to guarantee that revocation necessarily erases information about the plaintext. Broad-

[11] We remark that, there do exist post-quantum-insecure iO schemes based on well-founded assumptions [JLS21].
[12] Acording to the terminology of [AL21], this refers to finite term secure software leasing.

bent and Islam [BI20b] introduced the notion of *certified deletion*[13] and constructed a private-key quantum encryption scheme with the aforementioned feature which is inspired by the quantum key distribution protocol [BB84, TL17]. In contrast with Unruh's [Unr13] notion of revocable quantum ciphertexts which are eventually returned and verified, Broadbent and Islam [BI20b] consider certificates which are entirely classical. Moreover, the security definition requires that, once the certificate is successfully verified, the plaintext remains hidden even if the secret key is later revealed.

Using a hybrid encryption scheme, Hiroka, Morimae, Nishimaki and Yamakawa [HMNY21c] extended the scheme in [BI20a] to both public-key and attribute-based encryption with certified deletion via *receiver non-committing* encryption [JL00, CFGN96]. As a complementary result, the authors also gave a public-key encryption scheme with certified deletion which is *publicly verifiable* assuming the existence of one-shot signatures and extractable witness encryption. Using *Gaussian superpositions*, Poremba [Por22] proposed *Dual-Regev*-based public-key and fully homomorphic encryption schemes with certified deletion which are publicly verifiable and proven secure assuming a *strong Gaussian-collapsing conjecture* — a strengthening of the collapsing property of the Ajtai hash. Bartusek and Khurana [BK22] revisited the notion of certified deletion and presented a unified approach for how to generically convert any public-key, attribute-based, fully-homomorphic, timed-release or witness encryption scheme into an equivalent quantum encryption scheme with certified deletion. In particular, they considered a stronger notion called *certified everlasting security* which allows the adversary to be computationally unbounded once a valid deletion certificate is submitted.

Acknowledgements

We thank Fatih Kaleoglu and Ryo Nishimaki for several insightful discussions.

This work was done (in part) while the authors were visiting the Simons Institute for the Theory of Computing. P.A. is supported by a research gift from Cisco. A.P. is partially supported by AFOSR YIP (award number FA9550-16-1-0495), the Institute for Quantum Information and Matter (an NSF Physics Frontiers Center; NSF Grant PHY-1733907) and by a grant from the Simons Foundation (828076, TV). V.V. is supported by DARPA under Agreement No. HR00112020023, NSF CNS-2154149 and a Thornton Family Faculty Research Innovation Fellowship.

2 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.

2.1 Quantum Computing

For a comprehensive background on quantum computation, we refer to [NC11, Wil13]. 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

[13] This notion is incomparable with another related notion called unclonable encryption [BL20, AK21, AKL+22], which informally guarantees that it should be infeasible to clone quantum ciphertexts without losing information about the encrypted message.

by $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. 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 $\mathcal{H} = \left(\mathbb{C}^2\right)^{\otimes n}$ as n-qubit states. More generally, we associate *qudits* of dimension $d \geq 2$ with a d-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^d$. For brevity, we write $\mathcal{H}_d^n = \mathcal{H}_d^{\otimes n}$, where $\mathcal{H}_d$ is d-dimensional. 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 $\mathcal{H}$. Occasionally, we consider *subnormalized states*, i.e. states in the space of positive semidefinite operators over $\mathcal{H}$ with trace norm not exceeding 1.

The *trace distance* of two density matrices $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ is given by

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

Let $q \geq 2$ be a modulus and $n \in \mathbb{N}$ and let $\omega_q = e^{\frac{2\pi i}{q}} \in \mathbb{C}$ denote the primitive q-th root of unity. The n-qudit *q-ary quantum Fourier transform* over the ring $\mathbb{Z}_q^n$ is defined by the operation,

$$\mathsf{FT}_q : \quad |\mathbf{x}\rangle \quad \mapsto \quad \sqrt{q^{-n}} \sum_{\mathbf{y} \in \mathbb{Z}_q^n} \omega_q^{\langle \mathbf{x}, \mathbf{y} \rangle} |\mathbf{y}\rangle, \qquad \forall \mathbf{x} \in \mathbb{Z}_q^n.$$

The q-ary quantum Fourier transform is *unitary* and can be efficiently performed on a quantum computer for any modulus $q \geq 2$ [HH00].

A quantum channel $\Phi : \mathcal{L}(\mathcal{H}_A) \to \mathcal{L}(\mathcal{H}_B)$ is a linear map between linear operators over the Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$. Oftentimes, we use the compact notation $\Phi_{A \to B}$ 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 $\mathrm{Tr}[\Phi(X)] = \mathrm{Tr}[X]$, 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 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. Occasionally, we 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.

A polynomial-time *non-uniform* quantum algorithm is a family $\{(C_\lambda, \nu_\lambda)\}_{\lambda \in \mathbb{N}}$, where $\{C_\lambda\}_{\lambda \in \mathbb{N}}$ is a polynomial-size (not necessarily uniformly generated) family of circuits where, for each $\lambda \in \mathbb{N}$, a subset of input qubits to C_λ consists of a polynomial-size auxiliary density matrix ν_λ. We use the following notion of indistinguishability of quantum states in the presence of auxiliary inputs.

Definition 2.1 (Indistinguishability of ensembles of quantum states, [Wat05])**.** *Let $p : \mathbb{N} \to \mathbb{N}$ be a polynomially bounded function, and let ρ_λ and σ_λ be $p(\lambda)$-qubit quantum states. We say that $\{\rho_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{\sigma_\lambda\}_{\lambda \in \mathbb{N}}$ are quantum computationally indistinguishable ensembles of quantum states, denoted by $\rho_\lambda \approx_c \sigma_\lambda$, if, for any QPT distinguisher $\mathcal{D}$ with single-bit output, any polynomially bounded $q : \mathbb{N} \to \mathbb{N}$, any family of $q(\lambda)$-qubit auxiliary states $\{\nu_\lambda\}_{\lambda \in \mathbb{N}}$, and every $\lambda \in \mathbb{N}$,*

$$\left|\Pr[\mathcal{D}(\rho_\lambda \otimes \nu_\lambda) = 1] - \Pr[\mathcal{D}(\sigma_\lambda \otimes \nu_\lambda) = 1]\right| \leq \mathsf{negl}(\lambda).$$

Lemma 2.2 ("Almost As Good As New" Lemma, [Aar16]). *Let $\rho \in \mathcal{D}(\mathcal{H})$ be a density matrix over a Hilbert space $\mathcal{H}$. Let U be an arbitrary unitary and let $(\mathbf{\Pi}_0, \mathbf{\Pi}_1 = \mathbf{I} - \mathbf{\Pi}_0)$ be projectors acting on $\mathcal{H} \otimes \mathcal{H}_{\mathsf{aux}}$. We interpret $(U, \mathbf{\Pi}_0, \mathbf{\Pi}_1)$ as a measurement performed by appending an ancillary system in the state $|0\rangle\langle 0|_{\mathsf{aux}}$, applying the unitary U and subsequently performing the two-outcome measurement $\{\mathbf{\Pi}_0, \mathbf{\Pi}_1\}$ on the larger system. Suppose that the outcome corresponding to $\mathbf{\Pi}_0$ occurs with probability $1 - \varepsilon$, for some $\varepsilon \in [0, 1]$. In other words, it holds that $\mathrm{Tr}[\mathbf{\Pi}_0(U\rho \otimes |0\rangle\langle 0|_{\mathsf{aux}} U^\dagger)] = 1 - \varepsilon$. Then,*

$$\mathrm{TD}(\rho, \widetilde{\rho}) \leq \sqrt{\varepsilon},$$

where $\widetilde{\rho}$ is the state after performing the measurement and applying $U^\dagger$, and after tracing out $\mathcal{H}_{\mathsf{aux}}$:

$$\widetilde{\rho} = \mathrm{Tr}_{\mathsf{aux}}\left[U^\dagger \left(\mathbf{\Pi}_0 U(\rho \otimes |0\rangle\langle 0|_{\mathsf{aux}})U^\dagger \mathbf{\Pi}_0 + \mathbf{\Pi}_1 U(\rho \otimes |0\rangle\langle 0|_{\mathsf{aux}})U^\dagger \mathbf{\Pi}_1\right) U\right].$$

Lemma 2.3 (Quantum Union Bound, [Gao15]). *Let $\rho \in \mathcal{D}(\mathcal{H})$ be a state and let $\mathbf{\Pi}_1, \ldots, \mathbf{\Pi}_n \geq 0$ be sequence of (orthogonal) projections acting on $\mathcal{H}$. Suppose that, for every $i \in [n]$, it holds that $\mathrm{Tr}[\mathbf{\Pi}_i \rho] = 1 - \varepsilon_i$, for $\varepsilon_i \in [0, 1]$. Then, if we sequentially measure ρ with projective measurements $\{\mathbf{\Pi}_1, \mathbf{I} - \mathbf{\Pi}_1\}, \ldots, \{\mathbf{\Pi}_n, \mathbf{I} - \mathbf{\Pi}_n\}$, the probability that all measurements succeed is at least*

$$\mathrm{Tr}[\mathbf{\Pi}_n \cdots \mathbf{\Pi}_1 \rho \mathbf{\Pi}_1 \cdots \mathbf{\Pi}_n] \geq 1 - 4\sum_{i=1}^{n} \varepsilon_i.$$

2.2 Lattices and Cryptography

Let $n, m, p, q \in \mathbb{N}$ be positive integers. The rounding operation for $q \geq p \geq 2$ is the function

$$\lfloor \cdot \rceil_p \ : \ \mathbb{Z}_q \to \mathbb{Z}_p \ : \ x \mapsto \lfloor (p/q) \cdot x \rceil \pmod{p}.$$

A *lattice* $\Lambda \subset \mathbb{R}^m$ is a discrete subgroup of $\mathbb{R}^m$. Given a lattice $\Lambda \subset \mathbb{R}^m$ and a vector $\mathbf{t} \in \mathbb{R}^m$, we define the coset with respect to vector $\mathbf{t}$ as the lattice shift $\Lambda - \mathbf{t} = \{\mathbf{x} \in \mathbb{R}^m : \mathbf{x} + \mathbf{t} \in \Lambda\}$. Note that many different shifts $\mathbf{t}$ can define the same coset. The *dual* of a lattice $\Lambda \subset \mathbb{R}^m$, denoted by Λ^*, is the lattice of all $y \in \mathbb{R}^m$ that satisfy $\langle \mathbf{y}, \mathbf{x} \rangle \in \mathbb{Z}$, for every $\mathbf{x} \in \Lambda$. In other words, we let

$$\Lambda^* = \{\mathbf{y} \in \mathbb{R}^m \ : \ \langle \mathbf{y}, \mathbf{x} \rangle \in \mathbb{Z}, \ \text{for all } \mathbf{x} \in \Lambda\}.$$

In this work, we mainly consider *q-ary lattices* Λ that that satisfy $q\mathbb{Z}^m \subseteq \Lambda \subseteq \mathbb{Z}^m$, for some integer modulus $q \geq 2$. Specifically, we consider the lattice generated by a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ for some $n, m \in \mathbb{N}$ that consists of all vectors which are perpendicular to the rows of $\mathbf{A}$, namely

$$\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{0} \pmod{q}\}.$$

For any *syndrome* $\mathbf{y} \in \mathbb{Z}_q^n$ in the column span of $\mathbf{A}$, we also consider the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ given by

$$\Lambda_q^{\mathbf{y}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod{q}\} = \Lambda_q^\perp(\mathbf{A}) + \mathbf{c},$$

where $\mathbf{c} \in \mathbb{Z}^m$ is an arbitrary integer solution to the equation $\mathbf{A}\mathbf{c} = \mathbf{y} \pmod{q}$.

We use the following facts due to Gentry, Peikert and Vaikuntanathan [GPV07].

Lemma 2.4 ([GPV07], Lemma 5.1). *Let $n \in \mathbb{N}$ and let $q \geq 2$ be a prime modulus with $m \geq 2n \log q$. Then, for all but a q^{-n} fraction of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, the subset-sums of the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$. In other words, a uniformly random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ is full-rank with overwhelming probability.*

Gaussian Distribution. The *Gaussian measure* ρ_σ with parameter $\sigma > 0$ is defined as

$$\rho_\sigma(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/\sigma^2), \qquad \forall \mathbf{x} \in \mathbb{R}^m.$$

Let $\Lambda \subset \mathbb{R}^m$ be a lattice and let $\mathbf{t} \in \mathbb{R}^m$. We define the *Gaussian mass* of $\Lambda - \mathbf{t}$ as the quantity

$$\rho_\sigma(\Lambda - \mathbf{t}) = \sum_{\mathbf{y} \in \Lambda} \rho_\sigma(\mathbf{y} - \mathbf{t}).$$

The *discrete Gaussian distribution* $D_{\Lambda-\mathbf{t},\sigma}$ assigns probability proportional to $e^{-\pi\|\mathbf{x}\|^2/\sigma^2}$ to every vector $\mathbf{x} \in \Lambda - \mathbf{t}$. In other words, we have

$$D_{\Lambda-\mathbf{t},\sigma}(\mathbf{x}) = \frac{\rho_\sigma(\mathbf{x})}{\rho_\sigma(\Lambda - \mathbf{t})}, \qquad \forall \mathbf{x} \in \Lambda - \mathbf{t}.$$

In particular, for any coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ with $\mathbf{y} \in \mathbb{Z}_q^n$, the discrete Gaussian $D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}$ (centered around the origin) assigns probability proportional to $e^{-\pi\|\mathbf{x}\|^2/\sigma^2}$ to every vector $\mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A})$, and 0 otherwise.

The following lemma follows from [PR06, Lemma 2.11] and [GPV07, Lemma 5.3].

Lemma 2.5. *Let $n \in \mathbb{N}$ and let q be a prime with $m \geq 2n \log q$. Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$. Let $\mathbf{y} \in \mathbb{Z}_q^n$ be arbitrary. Then, for any $\sigma \geq \omega(\sqrt{\log m})$, there exists a negligible function $\varepsilon(m)$ such that*

$$D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}(\mathbf{x}) \leq 2^{-m} \cdot \frac{1+\varepsilon}{1-\varepsilon}, \qquad \forall\, \mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A}).$$

Let $\mathcal{B}^m(\mathbf{0}, r) = \{\mathbf{x} \in \mathbb{R}^m : \|\mathbf{x}\| \leq r\}$ denote the m-dimensional ball of radius $r > 0$. We use of the following tail bound for the Gaussian mass of a lattice [Ban93, Lemma 1.5 (ii)].

Lemma 2.6. *For any m-dimensional lattice Λ, shift $\mathbf{t} \in \mathbb{R}^m$, $\sigma > 0$ and $c \geq (2\pi)^{-\frac{1}{2}}$ it holds that*

$$\rho_\sigma\left((\Lambda - \mathbf{t}) \setminus \mathcal{B}^m(\mathbf{0}, c\sqrt{m}\sigma)\right) \leq (2\pi e c^2)^{\frac{m}{2}} e^{-\pi c^2 m} \rho_\sigma(\Lambda).$$

In addition, we also make use of the following tail bound for the discrete Gaussian which follows from [MR04, Lemma 4.4] and [GPV07, Lemma 5.3].

Lemma 2.7. *Let $n \in \mathbb{N}$ and let q be a prime with $m \geq 2n \log q$. Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$. Let $\mathbf{y} \in \mathbb{Z}_q^n$ be arbitrary. Then, for any $\sigma \geq \omega(\sqrt{\log m})$, there exists a negligible function $\varepsilon(m)$ such that*

$$\Pr_{\mathbf{x} \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}}\left[\|\mathbf{x}\| > \sigma\sqrt{m}\right] \leq 2^{-m} \cdot \frac{1+\varepsilon}{1-\varepsilon}.$$

A consequence of Lemma 2.6 is that the Gaussian distribution $D_{\mathbb{Z}^m,\sigma}$ is essentially only supported on the finite set $\{\mathbf{x} \in \mathbb{Z}^m : \|\mathbf{x}\| \leq \sigma\sqrt{m}\}$, which suggests the use of *truncation*.

Definition 2.8 (Truncated discrete Gaussian distribution)**.** *Let $m \in \mathbb{N}$, $q \geq 2$ be an integer modulus and let $\sigma > 0$ be a parameter. Then, the* truncated *discrete Gaussian distribution $D_{\mathbb{Z}_q^m,\sigma}$ with finite support $\{\mathbf{x} \in \mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m : \|\mathbf{x}\| \leq \sigma\sqrt{m}\}$ is defined as the density*

$$D_{\mathbb{Z}_q^m,\sigma}(\mathbf{x}) = \frac{\rho_\sigma(\mathbf{x})}{\sum_{\mathbf{y} \in \mathbb{Z}_q^m, \|\mathbf{y}\| \leq \sigma\sqrt{m}} \rho_\sigma(\mathbf{y})}.$$

Finally, we recall the following *noise smudging* property.

Lemma 2.9 (Noise smudging, [DGT+10])**.** *Let $y, \sigma > 0$. Then, the statistical distance between the distribution $D_{\mathbb{Z},\sigma}$ and $D_{\mathbb{Z},\sigma+y}$ is at most y/σ.*

We use the following technical lemma on the min-entropy of the truncated discrete Gaussian distribution, which we prove below.

Lemma 2.10. *Let $n \in \mathbb{N}$ and let q be a prime with $m \geq 2n \log q$. Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$. Then, for any $\sigma \geq \omega(\sqrt{\log m})$, there exists a negligible $\varepsilon(m)$ such that*

$$\max_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\substack{\mathbf{x} \in \mathbb{Z}_q^m, \|\mathbf{x}\| \leq \sigma\sqrt{m} \\ \mathbf{Ax} = \mathbf{y} \pmod q}} \left\{ \frac{\rho_\sigma(\mathbf{x})}{\sum_{\substack{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \leq \sigma\sqrt{m} \\ \mathbf{Az} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{z})} \right\} \leq 2^{-m+1} \cdot \frac{1+\varepsilon}{1-\varepsilon}.$$

Proof. Suppose that $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is a matrix whose columns generate $\mathbb{Z}_q^n$, i.e., $\mathbf{A}$ is full-rank. Then,

$$\begin{aligned}
&\max_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\substack{\mathbf{x} \in \mathbb{Z}_q^m, \|\mathbf{x}\| \leq \sigma\sqrt{m} \\ \mathbf{Ax} = \mathbf{y} \pmod q}} \left\{ \frac{\rho_\sigma(\mathbf{x})}{\sum_{\substack{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \leq \sigma\sqrt{m} \\ \mathbf{Az} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{z})} \right\} \\
&\leq \max_{\mathbf{y} \in \mathbb{Z}_q^n} \sup_{\mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A})} D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}(\mathbf{x}) \\
&\quad + \max_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\substack{\mathbf{x} \in \mathbb{Z}_q^m, \|\mathbf{x}\| \leq \sigma\sqrt{m} \\ \mathbf{Ax} = \mathbf{y} \pmod q}} \left| \frac{\rho_\sigma(\mathbf{x})}{\sum_{\substack{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \leq \sigma\sqrt{m} \\ \mathbf{Az} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{z})} - \frac{\rho_\sigma(\mathbf{x})}{\sum_{\substack{\mathbf{z} \in \mathbb{Z}^m \\ \mathbf{Az} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{z})} \right| \\
&\leq \max_{\mathbf{y} \in \mathbb{Z}_q^n} \sup_{\mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A})} D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}(\mathbf{x}) \\
&\quad + \max_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\substack{\mathbf{x} \in \mathbb{Z}_q^m, \|\mathbf{x}\| \leq \sigma\sqrt{m} \\ \mathbf{Ax} = \mathbf{y} \pmod q}} \frac{\rho_\sigma(\mathbf{x})}{\sum_{\substack{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \leq \sigma\sqrt{m} \\ \mathbf{Az} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{z})} \cdot \frac{\rho_\sigma(\Lambda_q^{\mathbf{y}}(\mathbf{A}) \setminus \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m}))}{\rho_\sigma(\Lambda_q^{\mathbf{y}}(\mathbf{A}))}
\end{aligned}$$

where $B^m(\mathbf{0}, r) = \{\mathbf{x} \in \mathbb{R}^m : \|\mathbf{x}\| \leq r\}$. Using the fact that

$$\frac{\rho_\sigma(\mathbf{x})}{\sum_{\substack{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \leq \sigma\sqrt{m} \\ \mathbf{Az} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{z})} \leq 1,$$

for $\mathbf{x} \in \mathbb{Z}_q^m$ with $\mathbf{Ax} = \mathbf{y} \pmod q$, and the fact that

$$\Pr_{\mathbf{v} \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}} \left[\|\mathbf{v}\| > \sigma\sqrt{m} \right] = \frac{\rho_\sigma(\Lambda_q^{\mathbf{y}}(\mathbf{A}) \setminus \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m}))}{\rho_\sigma(\Lambda_q^{\mathbf{y}}(\mathbf{A}))}$$

we get that

$$\max_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\substack{\mathbf{x} \in \mathbb{Z}_q^m, \|\mathbf{x}\| \leq \sigma\sqrt{m} \\ \mathbf{Ax} = \mathbf{y} \pmod q}} \left\{ \frac{\rho_\sigma(\mathbf{x})}{\sum_{\substack{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \leq \sigma\sqrt{m} \\ \mathbf{Az} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{z})} \right\}$$

$$\leq \max_{\mathbf{y} \in \mathbb{Z}_q^n} \left\{ \sup_{\mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A})} D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}(\mathbf{x}) + \Pr_{\mathbf{v} \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}} \left[\|\mathbf{v}\| > \sigma\sqrt{m} \right] \right\}.$$

Because $\sigma \geq \omega(\sqrt{\log m})$, the claim then follows from Lemma 2.5 and Lemma 2.7.

□

The Short Integer Solution problem. The *Short Integer Solution* (SIS) problem was introduced by Ajtai [Ajt96] in his seminal work on average-case lattice problems.

Definition 2.11 (Short Integer Solution problem, [Ajt96])**.** *Let $n, m \in \mathbb{N}$, $q \geq 2$ be a modulus and let $\beta > 0$ be a parameter. The Short Integer Solution problem ($\mathsf{SIS}^m_{n,q,\beta}$) problem is to find a short solution $\mathbf{x} \in \mathbb{Z}^m$ with $\|\mathbf{x}\| \leq \beta$ such that $\mathbf{A} \cdot \mathbf{x} = \mathbf{0} \pmod q$ given as input a matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$.*

Micciancio and Regev [MR07] showed that the SIS problem is, on the average, as hard as approximating worst-case lattice problems to within small factors. Subsequently, Gentry, Peikert and Vaikuntanathan [GPV07] gave an improved reduction showing that, for parameters $m = \text{poly}(n)$, $\beta = \text{poly}(n)$ and prime $q \geq \beta \cdot \omega(\sqrt{n \log q})$, the average-case $\mathsf{SIS}^m_{n,q,\beta}$ problem is as hard as approximating the shortest independent vector problem (SIVP) problem in the worst case to within a factor $\gamma = \beta \cdot \tilde{O}(\sqrt{n})$. We assume that $\mathsf{SIS}^m_{n,q,\beta}$, for $m = \Omega(n \log q)$, $\beta = 2^{o(n)}$ and $q = 2^{o(n)}$, is hard against quantum adversaries running in time $\text{poly}(q)$ with success probability $\text{poly}(1/q)$.

The Learning with Errors problem. The *Learning with Errors* problem was introduced by Regev [Reg05] and serves as the primary basis of hardness of post-quantum cryptosystems. The problem is defined as follows.

Definition 2.12 (Learning with Errors problem, [Reg05])**.** *Let $n, m \in \mathbb{N}$ be integers, let $q \geq 2$ be a modulus and let $\alpha \in (0, 1)$ be a noise ratio parameter. The (decisional) Learning with Errors ($\mathsf{LWE}^m_{n,q,\alpha q}$) problem is to distinguish between the following samples*

$$(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q) \quad \text{and} \quad (\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m),$$

where $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ is a uniformly random vector and where $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$ is a discrete Gaussian error vector. We rely on the quantum $\mathsf{LWE}^m_{n,q,\alpha q}$ assumption which states that the samples above are computationally indistinguishable for any QPT algorithm.

As shown in [Reg05], the $\mathsf{LWE}^m_{n,q,\alpha q}$ problem with parameter $\alpha q \geq 2\sqrt{n}$ is at least as hard as approximating the shortest independent vector problem (SIVP) to within a factor of $\gamma = \widetilde{O}(n/\alpha)$ in worst case lattices of dimension n. In this work we assume the subexponential hardness of $\mathsf{LWE}^m_{n,q,\alpha q}$ which relies on the worst case hardness of approximating short vector problems in lattices to within a subexponential factor. We assume that the $\mathsf{LWE}^m_{n,q,\alpha q}$ problem, for $m = \Omega(n \log q)$, $q = 2^{o(n)}$, $\alpha = 1/2^{o(n)}$, is hard against quantum adversaries running in time $\mathrm{poly}(q)$. We note that this parameter regime implies $\mathsf{SIS}^m_{n,q,\beta}$ [SSTX09].

Trapdoors for lattices. We use the following *trapdoor* property for the LWE problem.

Theorem 2.13 ([MP11], Theorem 5.1)**.** *Let $n, m \in \mathbb{N}$ and $q \in \mathbb{N}$ be a prime with $m = \Omega(n \log q)$. There exists a randomized algorithms with the following properties:*

- $\mathsf{GenTrap}(1^n, 1^m, q)$*: on input $1^n, 1^m$ and q, returns a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a trapdoor $\mathsf{td}_{\mathbf{A}}$ such that the distribution of $\mathbf{A}$ is negligibly (in the parameter n) close to uniform.*
- $\mathsf{Invert}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{b})$*: on input $\mathbf{A}$, $\mathsf{td}_{\mathbf{A}}$ and $\mathbf{b} = \mathbf{s}^\intercal \cdot \mathbf{A} + \mathbf{e}^\intercal \pmod q$, where $\|\mathbf{e}\| \leq q/(C_T\sqrt{n \log q})$ and $C_T > 0$ is a universal constant, returns $\mathbf{s}$ and $\mathbf{e}$ with overwhelming probability over $(\mathbf{A}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$.*

3 Quantum Discrete Gaussian Sampling for q-ary Lattices

In this section, we review some basic facts about Gaussian superpositions and present our *quantum discrete Gaussian sampler* which is used to revoke the decryption keys for our schemes.

3.1 Gaussian Superpositions

In this section, we review some basic facts about *Gaussian superpositions*. Given $q \in \mathbb{N}$, $m \in \mathbb{N}$ and $\sqrt{8m} < \sigma < q/\sqrt{8m}$, we consider Gaussian superpositions over $\mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$ of the form

$$|\psi\rangle = \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \rho_\sigma(\mathbf{x}) |\mathbf{x}\rangle .$$

Note that the state $|\psi\rangle$ is not normalized for convenience and ease of notation. The tail bound in Lemma 2.6 implies that (the normalized variant of) $|\psi\rangle$ is within negligible trace distance of a *truncated* discrete Gaussian superposition $|\tilde{\psi}\rangle$ with support $\{\mathbf{x} \in \mathbb{Z}_q^m : \|\mathbf{x}\| \leq \sigma\sqrt{\frac{m}{2}}\}$, where

$$|\tilde{\psi}\rangle = \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}(\mathbf{x})}\, |\mathbf{x}\rangle = \left(\sum_{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \leq \sigma\sqrt{\frac{m}{2}}} \rho_{\frac{\sigma}{\sqrt{2}}}(\mathbf{z}) \right)^{-\frac{1}{2}} \sum_{\mathbf{x} \in \mathbb{Z}_q^m : \|\mathbf{x}\| \leq \sigma\sqrt{\frac{m}{2}}} \rho_\sigma(\mathbf{x}) |\mathbf{x}\rangle .$$

In this work, we consider Gaussian superpositions with parameter $\sigma = \Omega(\sqrt{m})$ which can be efficiently implemented using standard quantum state preparation techniques; for example using *quantum rejection sampling* and the *Grover-Rudolph algorithm* [GR02, Reg05, Bra18, BCM+21].

Gaussian coset states. Our key-revocable encryption schemes in Section 6 and Section 8 rely on Gaussian superpositions over $\mathbf{x} \in \mathbb{Z}_q^m$ subject to a constraint of the form $\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod q$, for some matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and image $\mathbf{y} \in \mathbb{Z}_q^n$. In Algorithm 1, we give a procedure called $\mathsf{GenGauss}$ that, on input $\mathbf{A}$ and $\sigma > 0$, generates a Gaussian superposition state of the form

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m : \\ \mathbf{A}\mathbf{x} = \mathbf{y}}} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \, ,$$

for some $\mathbf{y} \in \mathbb{Z}_q^n$ which is statistically close to uniform whenever $m \geq 2n \log q$ and $\sigma \geq \omega(\sqrt{\log m})$. Because $|\psi_{\mathbf{y}}\rangle$ corresponds to a (truncated) Gaussian superposition over a particular lattice coset,

$$\Lambda_q^{\mathbf{y}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod q\},$$

of the q-ary lattice $\Lambda_q^{\perp}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{0} \pmod q\}$, we refer to it as a *Gaussian coset state*.

Finally, we recall an important property of Gaussian coset states.

Gaussian-collapsing hash functions. Unruh [Unr15] introduced the notion of collapsing hash functions in his seminal work on computationally binding quantum commitments. Informally, a hash function is called *collapsing* if it is computationally difficult to distinguish between a superposition of pre-images and a single (measured) pre-image.

In recent work, Poremba [Por22] proposed a special variant of the collapsing property with respect to *Gaussian superpositions*. Previously, Liu and Zhandry [LZ19] implicitly showed that the *Ajtai* hash function $h_{\mathbf{A}}(\mathbf{x}) = \mathbf{A}\mathbf{x} \pmod q$ is collapsing – and thus *Gaussian-collapsing* – via the notion of *lossy functions* and by assuming the superpolynomial hardness of (decisional) LWE.

We use the following result on the Gaussian-collapsing property of the Ajtai hash function.

Theorem 3.1 (Gaussian-collapsing property, [Por22], Theorem 4)**.** *Let $n \in \mathbb{N}$ and q be a prime with $m \geq 2n \log q$, each parameterized by $\lambda \in \mathbb{N}$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$. Then, the following samples are computationally indistinguishable assuming the quantum hardness of decisional $\mathsf{LWE}_{n,q,\alpha q}^m$, for any noise ratio $\alpha \in (0,1)$ with relative noise magnitude $1/\alpha = \sigma \cdot 2^{o(n)}$:*

$$\left(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \ |\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} = \mathbf{y}}} \rho_\sigma(\mathbf{x}) \ |\mathbf{x}\rangle \, , \ \mathbf{y} \in \mathbb{Z}_q^n \right) \approx_c \left(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \ |\mathbf{x}_0\rangle \, , \ \mathbf{A} \cdot \mathbf{x}_0 \in \mathbb{Z}_q^n \right)$$

where $(|\psi_{\mathbf{y}}\rangle \, , \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ and where $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ is a discrete Gaussian error.

3.2 Algorithm: GenGauss

The state preparation procedure $\mathsf{GenGauss}(\mathbf{A}, \sigma)$ is defined as follows.

Algorithm 1: $\mathsf{GenGauss}(\mathbf{A}, \sigma)$

Input: Matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and parameter $\sigma = \Omega(\sqrt{m})$.

Output: Gaussian state $|\psi_{\mathbf{y}}\rangle$ and $\mathbf{y} \in \mathbb{Z}_q^n$.

1 Prepare a Gaussian superposition in system X with parameter $\sigma > 0$:

$$|\psi\rangle_{XY} = \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle_X \otimes |\mathbf{0}\rangle_Y \, .$$

2 Apply the unitary $U_{\mathbf{A}} : |\mathbf{x}\rangle \, |\mathbf{0}\rangle \to |\mathbf{x}\rangle \, |\mathbf{A} \cdot \mathbf{x} \pmod q\rangle$ on systems X and Y:

$$|\psi\rangle_{XY} = \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle_X \otimes |\mathbf{A} \cdot \mathbf{x} \pmod q\rangle_Y \, .$$

3 Measure system Y in the computational basis, resulting in the state

$$|\psi_{\mathbf{y}}\rangle_{XY} = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m : \\ \mathbf{A}\mathbf{x} = \mathbf{y}}} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle_X \otimes |\mathbf{y}\rangle_Y \, .$$

4 Output the state $|\psi_{\mathbf{y}}\rangle$ in system X and the outcome $\mathbf{y} \in \mathbb{Z}_q^n$ in system Y.

3.3 Algorithm: QSampGauss

Recall that, in Algorithm 1, we gave a procedure called $\mathsf{GenGauss}(\mathbf{A}, \sigma)$ that prepares a Gaussian coset state $|\psi_{\mathbf{y}}\rangle$, for a randomly generated $\mathbf{y} \in \mathbb{Z}_q^n$. In general, however, generating a specific Gaussian coset state on input $(\mathbf{A}, \mathbf{y})$ requires a *short trapdoor basis* $\mathsf{td}_{\mathbf{A}}$ for the matrix $\mathbf{A}$. This task can be thought of as a quantum analogue of the *discrete Gaussian sampling problem* [GPV07], where the goal is to output a sample $\mathbf{x} \sim D_{Z^m, \sigma}$ such that $\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod q$ on input $(\mathbf{A}, \mathbf{y})$ and $\sigma > 0$.

In Algorithm 2, we give a procedure called $\mathsf{QSampGauss}$ which, on input $(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ generates a specific Gaussian coset state $|\psi_{\mathbf{y}}\rangle$ of the form

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} = \mathbf{y}}} \rho_\sigma(\mathbf{x}) \; |\mathbf{x}\rangle \, .$$

Our procedure $\mathsf{QSampGauss}$ in Algorithm 2 can be thought of as an explicit quantum reduction from $\mathsf{ISIS}^m_{n,q,\sigma\sqrt{m/2}}$ to $\mathsf{LWE}^m_{n,q,q/\sqrt{2}\sigma}$ which is inspired by the quantum reduction of Stehlé et al. [SSTX09] which reduces SIS to LWE. To obtain the aforementioned reduction, one simply needs to replace the procedure $\mathsf{Invert}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \cdot)$ in Step 4 in Algorithm 2 with a solver for the $\mathsf{LWE}^m_{n,q,q/\sqrt{2}\sigma}$ problem.

In Theorem 3.3, we prove the correctness of Algorithm 2. As a technical ingredient, we rely on a *duality lemma* [Por22] that characterizes the Fourier transform of a Gaussian coset state in terms

of its dual state. Note that $|\psi_{\mathbf{y}}\rangle$ corresponds to a Gaussian superposition over a lattice coset,

$$\Lambda_q^{\mathbf{y}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod q\},$$

of the q-ary lattice $\Lambda_q^{\perp}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{0} \pmod q\}$. Here, the *dual* of $\Lambda_q^{\perp}(\mathbf{A})$ satisfies $q \cdot \Lambda_q^{\perp}(\mathbf{A})^* = \Lambda_q(\mathbf{A})$, where $\Lambda_q(\mathbf{A})$ corresponds to the lattice generated by $\mathbf{A}^\intercal$, i.e.

$$\Lambda_q(\mathbf{A}) = \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{z} = \mathbf{A}^\intercal \cdot \mathbf{s} \pmod q, \text{ for some } \mathbf{s} \in \mathbb{Z}^n\}.$$

The following lemma relates the Fourier transform of $|\psi_{\mathbf{y}}\rangle$ with a superposition of LWE samples with respect to a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a phase which depends on $\mathbf{y}$. In other words, the resulting state can be thought of as a superposition of Gaussian balls around random lattice vectors in $\Lambda_q(\mathbf{A})$.

Lemma 3.2 ([Por22], Lemma 16)**.** *Let $m \in \mathbb{N}$, $q \geq 2$ be a prime and $\sqrt{8m} < \sigma < q/\sqrt{8m}$. Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$ and let $\mathbf{y} \in \mathbb{Z}_q^n$ be arbitrary. Then, the q-ary quantum Fourier transform of the (normalized variant of the) Gaussian coset state*

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x}) |\mathbf{x}\rangle$$

is within negligible (in $m \in \mathbb{N}$) trace distance of the (normalized variant of the) Gaussian state

$$|\hat{\psi}_{\mathbf{y}}\rangle = \sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \rho_{q/\sigma}(\mathbf{e})\, \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle} |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q\rangle\,.$$

The procedure $\mathsf{QSampGauss}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ is defined as follows.

Algorithm 2: $\mathsf{QSampGauss}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$

Input: Matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a trapdoor $\mathsf{td}_{\mathbf{A}}$, an image $\mathbf{y} \in \mathbb{Z}_q^n$ and parameter $\sigma = O(\frac{q}{\sqrt{m}})$.

Output: Gaussian state $|\psi_{\mathbf{y}}\rangle$.

1 Prepare the following superposition with parameter $q/\sigma > 0$:

$$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} |\mathbf{s}\rangle \otimes \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \rho_{q/\sigma}(\mathbf{e}) \, |\mathbf{e}\rangle \otimes |\mathbf{0}\rangle$$

2 Apply the generalized Pauli operator $\mathbf{Z}_q^{-\mathbf{y}}$ on the first register, resulting in the state

$$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle} \, |\mathbf{s}\rangle \otimes \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \rho_{q/\sigma}(\mathbf{e}) \, |\mathbf{e}\rangle \otimes |\mathbf{0}\rangle$$

3 Apply the unitary $U_{\mathbf{A}} : |\mathbf{s}\rangle \, |\mathbf{e}\rangle \, |\mathbf{0}\rangle \to |\mathbf{s}\rangle \, |\mathbf{e}\rangle \, |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q\rangle$, resulting in the state

$$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \rho_{q/\sigma}(\mathbf{e}) \, \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle} \, |\mathbf{s}\rangle \, |\mathbf{e}\rangle \, |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q\rangle$$

4 Coherently run $\mathsf{Invert}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \cdot)$ on the third register in order to uncompute the first and the second register, resulting in a state that is close in trace distance to the following state:

$$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \rho_{q/\sigma}(\mathbf{e}) \, \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle} \, |0\rangle \, |0\rangle \, |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q\rangle$$

5 Discard the first two registers. Apply the (inverse) quantum Fourier transform and output the resulting state.

Let us now prove the correctness of Algorithm 2.

Theorem 3.3 (Quantum Discrete Gaussian Sampler)**.** *Let $n \in \mathbb{N}$, q be a prime with $m \geq 2n \log q$ and $\sqrt{8m} < \sigma < q/\sqrt{8m}$. Let $(\mathbf{A}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ be sampled as in Theorem 2.13 and let $\mathbf{y} \in \mathbb{Z}_q^n$ be arbitrary. Then, with overwhelming probability, $\mathsf{QSampGauss}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ in Algorithm 2 outputs a state which is within negligible trace distance of the (normalized variant of the) state,*

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \, .$$

Proof. From Lemma 2.4 and Theorem 2.13, it follows that $(\mathbf{A}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ yields a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$ with overwhelming probability. Moreover, since $\sqrt{8m} < \sigma < q/\sqrt{8m}$, the inversion procedure $\mathsf{Invert}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \cdot)$ from Theorem 2.13

in Step 4 in Algorithm 2 succeeds with overwhelming probability at generating the Gaussian state

$$|\hat{\psi}_{\mathbf{y}}\rangle = \sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m} \rho_{q/\sigma}(\mathbf{e})\,\omega_q^{-\langle \mathbf{s},\mathbf{y}\rangle}\,|\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal \pmod q\rangle$$

Applying the (inverse) quantum Fourier transform $\mathsf{FT}_q^\dagger$, the claim then follows from Lemma 3.2. □

4 Quantum Goldreich-Levin Theorem for Large Fields

In this section, we give a proof of a quantum Goldreich-Levin theorem for large fields $\mathbb{Z}_q$.

4.1 Post-Quantum Reductions and Quantum Rewinding

We first review some recent work by Bitansky, Brakerski and Kalai [BBK22] that enables us to convert a wide range of classical reductions into post-quantum reductions (which allow for quantum auxiliary input) in a constructive manner. We first review some basic terminology from [BBK22].

Let $\lambda \in \mathbb{N}$ be a parameter. A *non-interactive assumption* $\mathsf{P} = (\mathsf{G}, \mathsf{V}, c)$ with respect to a set of polynomials $d(\lambda), n(\lambda)$ and $m(\lambda)$ is characterized as follows:

- The generator G takes as input 1^λ and $r \in \{0,1\}^d$, and returns $x \in \{0,1\}^n$.
- The verifier V takes as input 1^λ and $(r, y) \in \{0,1\}^d \times \{0,1\}^m$, and returns a single bit output.
- $c(\lambda)$ is the threshold associated with the assumption.

Given a (possibly randomized) *solver*, we characterize the *advantage* in solving an assumption P in terms of the absolute distance between the solving probability (or, *value*) and the threshold c; for example, for a *decision assumption* P (with $m = 1$) we characterize the value in solving P in terms of $\frac{1}{2} + \varepsilon$, where the threshold is given by $c(\lambda) = \frac{1}{2}$ and $\varepsilon > 0$ is corresponds to the *advantage*. We say that a reduction is *black-box* if it is oblivious to the representation and inner workings of the solver that is being used. Moreover, we say that a reduction is *non-adaptive* if all queries to the solver are known ahead of time.

We use the following theorem.

Theorem 4.1 ([BBK22], adapted from Theorem 7.1)**.** *Let $c \in \mathbb{R}$. Suppose that there exists a classical reduction from solving a non-interactive assumption Q to solving a non-interactive assumption P such that the following holds: if the P-solver has advantage $\varepsilon > 0$ then the Q-solver has advantage c (independent of ε) with running time $\mathrm{poly}(1/\varepsilon, c, \lambda)$.*

Then, there exists a quantum reduction from solving Q to quantumly solving P such that the following holds: if the quantum P-solver (with non-uniform quantum advice) has an advantage given by $\varepsilon > 0$, then the Q-solver has advantage c (the same as the classical reduction) with running time $\mathrm{poly}(1/\varepsilon, c, \lambda)$.

4.2 Goldreich-Levin Theorems for Large Fields

The following result is implicit in the work of Dodis et al. [DGT+10].

Theorem 4.2 (Classical Goldreich-Levin Theorem for Finite Fields, [DGT+10], Theorem 1)**.** *Let q be a prime and $m \in \mathbb{N}$. Let $H = \{\mathbf{x} \in \mathbb{Z}_q^m \,:\, \|\mathbf{x}\| \leq \sigma\sqrt{m}\}$ be a subset of $\mathbb{Z}_q^m$, for some $\sigma > 0$. Let $f : H \to \{0,1\}^*$ be any (possibly randomized) function. Suppose there exists a distinguisher $\mathcal{D}$ that runs in time $T(\mathcal{D})$ and has the property that*

$$\left|\Pr\left[\mathcal{D}(\mathbf{u}, \mathbf{u}^\intercal\mathbf{x}, \mathsf{aux}) = 1 \,:\, \begin{smallmatrix}\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m \\ \mathbf{x} \sim D_{\mathbb{Z}_q^m,\sigma} \\ \mathsf{aux} \leftarrow f(\mathbf{x})\end{smallmatrix}\right] - \Pr\left[\mathcal{D}(\mathbf{u}, r, \mathsf{aux}) = 1 \,:\, \begin{smallmatrix}\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m, r \xleftarrow{\$} \mathbb{Z}_q \\ \mathbf{x} \sim D_{\mathbb{Z}_q^m,\sigma} \\ \mathsf{aux} \leftarrow f(\mathbf{x})\end{smallmatrix}\right]\right| = \varepsilon.$$

Then, there exists a (classical) non-adaptive black-box extractor $\mathcal{E}$ whose running time is given by $T(\mathcal{E}) = T(\mathcal{D}) \cdot \mathrm{poly}(m, \sigma, 1/\varepsilon)$ and succeeds with probability at least

$$\Pr\left[\mathcal{E}(\mathsf{aux}) = \mathbf{x} \,:\, \mathsf{aux} \leftarrow f(\mathbf{x})\right] \geq \frac{\varepsilon^3}{512 \cdot m \cdot q^2}.$$

Using the constructive post-quantum reduction from Theorem 4.1, we can convert Theorem 4.2 into a quantum Goldreich-Levin Theorem for finite fields, and obtain the following.

Theorem 4.3 (Quantum Goldreich-Levin Theorem for Large Fields)**.** *Let q be a prime and $m \in \mathbb{N}$. Let $H = \{\mathbf{x} \in \mathbb{Z}_q^m \,:\, \|\mathbf{x}\| \leq \sigma\sqrt{m}\}$ be a subset of $\mathbb{Z}_q^m$, for some $\sigma > 0$. Let $\Phi : \mathcal{L}(\mathcal{H}_q^m) \to \mathcal{L}(\mathcal{H}_{\mathrm{AUX}})$ be any* CPTP *map with auxiliary system $\mathcal{H}_{\mathrm{AUX}}$. Suppose there exists a quantum distinguisher $\mathcal{D}$ that runs in time $T(\mathcal{D})$ and has the property that*

$$\left|\Pr\left[\mathcal{D}(\mathbf{u}, \mathbf{u}^\intercal\mathbf{x}, \mathsf{aux}) = 1 \,:\, \begin{smallmatrix}\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m \\ \mathbf{x} \sim D_{\mathbb{Z}_q^m,\sigma} \\ \mathsf{aux} \leftarrow \Phi(|\mathbf{x}\rangle\langle\mathbf{x}|)\end{smallmatrix}\right] - \Pr\left[\mathcal{D}(\mathbf{u}, r, \mathsf{aux}) = 1 \,:\, \begin{smallmatrix}\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m, r \xleftarrow{\$} \mathbb{Z}_q \\ \mathbf{x} \sim D_{\mathbb{Z}_q^m,\sigma} \\ \mathsf{aux} \leftarrow \Phi(|\mathbf{x}\rangle\langle\mathbf{x}|)\end{smallmatrix}\right]\right| = \varepsilon.$$

Then, there exists a quantum extractor $\mathcal{E}$ whose running time is given by $T(\mathcal{E}) = T(\mathcal{D}) \cdot \mathrm{poly}(m, \sigma, 1/\varepsilon)$ and that succeeds with probability at least

$$\Pr\left[\mathcal{E}(\mathsf{aux}) = \mathbf{x} \,:\, \mathsf{aux} \leftarrow \Phi(|\mathbf{x}\rangle\langle\mathbf{x}|)\right] \geq \mathrm{poly}(\varepsilon, 1/m, 1/\sigma, 1/q).$$

Proof. The proof follows immediately by combining Theorem 4.2 and Theorem 4.1. □

5 Definition: Key-Revocable Public-Key Encryption

Let us now give a formal definition of key-revocable public-key encryption schemes.

Definition 5.1 (Key-Revocable Public-Key Encryption)**.** *A key-revocable public-key encryption scheme consists of efficient algorithms* $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$*, where* Enc *is a* PPT *algorithm and* KeyGen, Dec *and* Revoke *are* QPT *algorithms defined as follows:*

- $\mathsf{KeyGen}(1^\lambda)$*: given as input a security parameter λ, output a public key* PK*, a master secret key* MSK *and a quantum decryption key ρ_{SK}.*
- $\mathsf{Enc}(\mathsf{PK}, x)$*: given a public key* PK *and plaintext $x \in \{0,1\}^\ell$, output a ciphertext* CT*.*
- $\mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT})$*: given a decryption key ρ_{SK} and ciphertext* CT*, output a message y.*
- $\mathsf{Revoke}(\mathsf{PK}, \mathsf{MSK}, \sigma)$*: given as input a master secret key* MSK*, a public key* PK *and quantum state σ, output* Valid *or* Invalid*.*

Correctness of Decryption. For every $x \in \{0,1\}^\ell$, the following holds:

$$\Pr\left[x \leftarrow \mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT}) \ : \ \begin{smallmatrix}(\mathsf{PK},\mathsf{MSK},\rho_{\mathsf{SK}})\leftarrow\mathsf{KeyGen}(1^\lambda)\\ \mathsf{CT}\leftarrow\mathsf{Enc}(\mathsf{PK},x)\end{smallmatrix}\right] \geq 1-\nu(\lambda),$$

where $\nu(\cdot)$ is a negligible function.

Correctness of Revocation. The following holds:

$$\Pr\left[\mathsf{Valid} \leftarrow \mathsf{Revoke}\,(\mathsf{PK},\mathsf{MSK},\rho_{\mathsf{SK}}) \ : \ (\mathsf{PK},\mathsf{MSK},\rho_{\mathsf{SK}}) \leftarrow \mathsf{KeyGen}(1^\lambda)\right] \geq 1-\nu(\lambda),$$

where $\nu(\cdot)$ is a negligible function.

Remark 5.2. *Using the well-known "Almost As Good As New Lemma" (Lemma 2.2), the procedure* Dec *can be purified to obtain another quantum circuit* $\widetilde{\mathsf{Dec}}$ *such that* $\widetilde{\mathsf{Dec}}(\rho_{\mathsf{SK}}, \mathsf{CT})$ *yields* (x, ρ'_{SK}) *with probability at least* $1-\nu(\lambda)$ *and moreover,* $\mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{PK}, x)$ *and* $\mathrm{TD}(\rho'_{\mathsf{SK}}, \rho_{\mathsf{SK}}) \leq \nu'(\lambda)$ *with* $\nu'(\lambda)$ *is another negligible function.*

5.1 Security Definition

Our security definition for key-revocable public-key encryption is as follows.

Definition 5.3. *A key-revocable public-key encryption scheme* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ *is* (ϵ,δ)*-secure if, for every* QPT *adversary* $\mathcal{A}$ *with*

$$\Pr[\mathsf{Invalid} \leftarrow \mathsf{Expt}_{\Sigma,\mathcal{A}}(1^\lambda, b)] \leq \delta(\lambda)$$

for $b \in \{0,1\}$*, it holds that*

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}_{\Sigma,\mathcal{A}}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}_{\Sigma,\mathcal{A}}(1^\lambda, 1)\right]\right| \leq \varepsilon(\lambda),$$

where $\mathsf{Expt}_{\Sigma,\mathcal{A}}(1^\lambda, b)$ *is as defined in Figure 1. If* $\delta(\lambda) = 1 - 1/\mathrm{poly}(\lambda)$ *and* $\varepsilon(\lambda) = \mathsf{negl}(\lambda)$*, we simply say the key-revocable public-key encryption scheme is secure.*

Remark 5.4. *Our security definition is similar to the one proposed by Agrawal et al. [AKN+23] in the context of public-key encryption with secure leasing.*

5.2 Key-Revocable Public-Key Fully Homomorphic Encryption

A key-revocable public-key fully homomorphic encryption scheme defined for a class of functions $\mathcal{F}$, in addition to $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$, consists of the following PPT algorithm:

- $\mathsf{Eval}(\mathsf{PK}, f, \mathsf{CT})$: on input a public key PK, function $f \in \mathcal{F}$, ciphertext CT, outputs another ciphertext CT'.

Remark 5.5. *Sometimes we allow* KeyGen *to additionally take as input different parameters associated with the implementations of the functions in* $\mathcal{F}$*. For example, we allow* KeyGen *to take as input a parameter* L *in such a way that all the parameters in the system depend on* L *and moreover, the homomorphic evaluation is only supported on circuits (in* $\mathcal{F}$*) of depth at most* L*.*

$\underline{\mathsf{Expt}_{\Sigma,\mathcal{A}}\left(1^\lambda, b\right)}$:

Initialization Phase:

- The challenger runs $(\mathsf{PK}, \mathsf{MSK}, \rho_{\mathsf{SK}}) \leftarrow \mathsf{KeyGen}(1^\lambda)$ and sends $(\mathsf{PK}, \rho_{\mathsf{SK}})$ to $\mathcal{A}$.

Revocation Phase:

- The challenger sends the message `REVOKE` to $\mathcal{A}$.
- The adversary $\mathcal{A}$ returns a state σ.
- The challenger aborts if $\mathsf{Revoke}(\mathsf{PK}, \mathsf{MSK}, \sigma)$ outputs $\mathsf{Invalid}$.

Guessing Phase:

- $\mathcal{A}$ submits a plaintext $x \in \{0,1\}^\ell$ to the challenger.
- If $b = 0$: The challenger sends $\mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{PK}, x)$ to $\mathcal{A}$. Else, if $b = 1$, the challenger sends $\mathsf{CT} \xleftarrow{\$} \mathcal{C}$, where $\mathcal{C}$ is the ciphertext space of ℓ bit messages.
- Output $b_{\mathcal{A}}$ if the output of $\mathcal{A}$ is $b_{\mathcal{A}}$.

Figure 1: Security Experiment

Correctness of Evaluation and Decryption. For every $f \in \mathcal{F}$ with ℓ-bit inputs, every $x \in \{0,1\}^\ell$, the following holds:

$$\mathsf{Pr}\left[f(x) \leftarrow \mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT}') \;:\; \begin{matrix}(\mathsf{PK},\mathsf{MSK},\rho_{\mathsf{SK}}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{PK}, x) \\ \mathsf{CT}' \leftarrow \mathsf{Eval}(\mathsf{PK}, f, \mathsf{CT})\end{matrix}\right] \geq 1 - \nu(\lambda),$$

where $\nu(\cdot)$ is a negligible function.

Correctness of Revocation. Defined as before.

Security. Defined as before (Definition 5.3).

6 Key-Revocable Dual-Regev Encryption

In this section, we present a construction of key-revocable public-key encryption from learnin with errors. Our construction essentially involves making the Dual Regev public-key encryption of Gentry, Peikert and Vaikuntanathan [GPV07] key revocable.

6.1 Construction

We define our Dual-Regev construction below.

Construction 1 (Key-Revocable Dual-Regev Encryption)**.** *Let $n \in \mathbb{N}$ be the security parameter and $m \in \mathbb{N}$. Let $q \geq 2$ be a prime and let $\alpha, \beta, \sigma > 0$ be parameters. The key-revocable public-key scheme* $\mathsf{RevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ *consists of the following* QPT *algorithms:*

- $\mathsf{KeyGen}(1^\lambda) \to (\mathsf{PK}, \rho_{\mathsf{SK}}, \mathsf{MSK})$: *sample* $(\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ *and generate a Gaussian superposition* $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ *with*

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} = \mathbf{y}}} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \,,$$

  for some $\mathbf{y} \in \mathbb{Z}_q^n$. Output $\mathsf{PK} = (\mathbf{A}, \mathbf{y})$, $\rho_{\mathsf{SK}} = |\psi_{\mathbf{y}}\rangle$ and $\mathsf{MSK} = \mathsf{td}_{\mathbf{A}}$.

- $\mathsf{Enc}(\mathsf{PK}, \mu) \to \mathsf{CT}$: *to encrypt a bit $\mu \in \{0,1\}$, sample a random vector $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ and errors $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$ and $e' \sim D_{\mathbb{Z}, \beta q}$, and output the ciphertext pair*

$$\mathsf{CT} = \left(\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q, \mathbf{s}^\intercal \mathbf{y} + e' + \mu \cdot \lfloor \frac{q}{2} \rfloor \pmod q\right) \in \mathbb{Z}_q^m \times \mathbb{Z}_q.$$

- $\mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT}) \to \{0,1\}$: *to decrypt* CT*, apply the unitary $U : |\mathbf{x}\rangle |0\rangle \to |\mathbf{x}\rangle |\mathsf{CT} \cdot (-\mathbf{x}, 1)^\intercal\rangle$ on input $|\psi_{\mathbf{y}}\rangle |0\rangle$, where $\rho_{\mathsf{SK}} = |\psi_{\mathbf{y}}\rangle$, and measure the second register in the computational basis. Output 0, if the measurement outcome is closer to 0 than to $\lfloor \frac{q}{2} \rfloor$, and output 1, otherwise.*

- $\mathsf{Revoke}(\mathsf{MSK}, \mathsf{PK}, \rho) \to \{\top, \bot\}$*: on input* $\mathsf{td}_{\mathbf{A}} \leftarrow \mathsf{MSK}$ *and* $(\mathbf{A}, \mathbf{y}) \leftarrow \mathsf{PK}$*, apply the measurement $\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, I - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\}$ onto the state ρ using the procedure* $\mathsf{QSampGauss}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ *in Algorithm 2. Output $\top$, if the measurement is successful, and output $\bot$ otherwise.*

Correctness of Decryption. Follows from the correctness of Dual-Regev public-key encryption.

Correctness of Revocation. This follows from Theorem 3.3.

Let us now prove the security of our key-revocable Dual-Regev scheme in Construction 1. Our first result concerns $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda))$-security, i.e., we assume that revocation succeeds with overwhelming probability.

Theorem 6.1. *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and let $\alpha, \beta \in (0,1)$ be noise ratios chosen such that $\beta/\alpha = 2^{o(n)}$ and $1/\alpha = 2^{o(n)} \cdot \sigma$. Then, assuming the subexponential hardness of the $\mathsf{LWE}^m_{n,q,\alpha q}$ and $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ problems, the scheme* $\mathsf{RevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ *in Construction 1 is a $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda))$-secure key-revocable public-key encryption scheme according to Definition 5.3.*

To prove the stronger variant of $(\mathsf{negl}(\lambda), 1 - 1/\mathrm{poly}(\lambda))$-security, i.e., where we do not make any requirements about the success probability of revocation, we need to invoke Theorem 6.16.

Theorem 6.2. *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and let $\alpha, \beta \in (0,1)$ be noise ratios chosen such that $\beta/\alpha = 2^{o(n)}$ and $1/\alpha = 2^{o(n)} \cdot \sigma$. Assuming Theorem 6.16, the scheme* $\mathsf{RevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ *in Construction 1 is a $(\mathsf{negl}(\lambda), 1 - 1/\mathrm{poly}(\lambda))$-secure key-revocable public-key encryption scheme according to Definition 5.3.*

Guide for proving Theorem 6.1 and Theorem 6.2.

- First, we prove a technical lemma (Lemma 6.10) that helps us remove the condition that revocation succeeds when analyzing the advantage of a distinguisher. Our proof uses *projective implementations* which allow us to estimate the success probability of quantum programs.
- The next step towards proving Theorem 6.1 is a search-to-decision reduction with quantum auxiliary input for the Dual-Regev scheme (Theorem 6.11). Here, we show how to extract a short vector mapping $\mathbf{A}$ to $\mathbf{y}$ from an efficient adversary who has a non-negligible distinguishing advantage at distinguishing Dual-Regev ciphertexts from uniform.
- Next, we state the *Simultaneous Dual-Regev Extraction* conjecture in Theorem 6.16, which is a strengthening of our search-to-decision reduction in Theorem 6.11. Informally, it says that extraction of a short vector mapping $\mathbf{A}$ to $\mathbf{y}$ succeeds, even if we apply revocation on a separate register. We prove that Theorem 6.16 holds assuming $\mathsf{LWE}/\mathsf{SIS}$ in the special case when revocation succeeds with overwhelming probability. This is captured by Theorem 6.19.
- Next, we prove technical lemma which exploits the search-to-reduction to extract two *distinct* short vectors mapping $\mathbf{A}$ to $\mathbf{y}$. This is proven in Section 6.4.
- Finally, we put all the pieces together in Section 6.5 and show how to use the result from Section 6.4 in order to break the SIS assumption.

6.2 Threshold Implementations

In this section, we prove Lemma 6.10. This is a useful ingredient in the security proofs behind our key-revocable Dual-Regev encryption scheme.

First, we review some recent techniques that allow us to measure the success probability of *quantum programs*. In the classical setting, this task is fairly straightforward: simply execute a given program on samples from a *test distribution*, and check how many times the program succeeds. Using standard concentration inequalities, one can then estimate the success probability to inverse polymonial precision. In the quantum realm, however, this task is non-trivial if the quantum program is run with respect to quantum auxiliary inputs.

Inspired by the work of Marriott and Watrous [MW05], Zhandry [Zha20] introduced the notion of projective implementations which allow us to accomplish this task efficiently. Below, we introduce some relevant definitions and results from the original work of Zhandry [Zha20], as well as subsequent follow-up works [ALL+21, CLLZ21, AKL+22]. First, we discuss *inefficient* measurement techniques for measuring the success probability of a quantum program. Next, we move onto *efficient* measurement techniques that allow us to obtain such estimates approximately.

Inefficient measurements. Suppose we we have quantum program, say consisting of a quantum circuit and some quantum auxiliary input, and we wish to estimate its success probability. A natural starting point is to consider a two-outcome POVM $\mathcal{P} = (P, Q)$ over the two outcomes 0 (success) and 1 (failure). Zhandry [Zha20] showed that for any such $\mathcal{P}$, there exists a natural projective measurement (called a *projective implementation*) such that the post-measurement state corresponds precisely to an eigenvector of P. Moreover, there exists a projective measurement $\mathcal{E}$ that *measures* the success probability with respect to $\mathcal{P}$ on some auxiliary input state; specifically,

- $\mathcal{E}$ outputs a probability $p \in [0,1]$ (i.e., a real number) from the set of eigenvalues of P.
- The post-measurement state after obtaining outcome p corresponds to an eigenvector of P with eigenvalue p; similarly, it is an eigenvector of $Q = I - P$ with eigenvalue $1-p$.

The measurement $\mathcal{E}$ is projective in the following sense: whenever we apply the same measurement $\mathcal{E}$ on the post-measurement state, we obtain precisely the same outcome. The following theorem is implicit in [Zha20, Lemma 1], but we rely on the presentation from [AKL+22, Theorem 2.5].

Theorem 6.3 (Projective implementation)**.** *Let $\mathcal{P} = (P,Q)$ be a two-outcome POVM and let $\mathcal{D}$ be the distribution over the eigenvalues of P. Then, there exists a projective measurement $\mathcal{E} = \{E_p\}_{p\in\mathcal{D}}$ with index set $\mathcal{D}$ such that: for every quantum state ρ, where we let $\rho_p = E_p \rho E_p$ denote the sub-normalized post-measurement state after measuring ρ via E_p, it holds that*

- *For every $p \in \mathcal{D}$, the state ρ_p is an eigenvector of P with eigenvaue p, and*
- *the probability of ρ when measured with respect to P is equal to $\mathrm{Tr}[P\rho] = \sum_{p\in\mathcal{D}} \mathrm{Tr}[P\rho_p]$.*

Remark 6.4. *Suppose that $\mathcal{P} = (P,Q)$ is a two-outcome POVM and that P has an eigenbasis $\{|\psi_i\rangle\}$ with associated eigenvalues $\{\lambda_i\}$. Because P and Q commute, they share a common eigenbasis. In this case, there exists a natural measurement $\mathcal{E}$ that corresponds to a projective implementation of the POVM $\mathcal{P}$; namely, for any input $|\psi\rangle$, which can express as $|\psi\rangle = \sum_i \alpha_i |\psi_i\rangle$, the measurement $\mathcal{E} = \{E_{\lambda_i}\}$ will result in outcome λ_i and a leftover eigenstate $|\psi_i\rangle$ with probabiliy $|\alpha_i|^2$.*

Next, we use a generalization of projective implementations introduced in [ALL+21]. Rather than estimating the success probability directly, we can instead measure whether it is above or below a certain threshold. This gives rise to the following notion of *threshold implementations*.

Theorem 6.5 (Threshold implementation)**.** *Let $\gamma \in (0,1)$ be a parameter and let $\mathcal{P} = (P,Q)$ be a two-outcome POVM, where P has an eigenbasis $\{|\psi_i\rangle\}$ with associated eigenvalues $\{\lambda_i\}$. Then, there exists a projective threshold implementation $(\mathsf{TI}_\gamma(\mathcal{P}), \mathrm{I} - \mathsf{TI}_\gamma(\mathcal{P}))$ such that*

- *$\mathsf{TI}_\gamma(\mathcal{P})$ projects a quantum state into the subspace spanned by $\{|\psi_i\rangle\}$ whose eigenvalues λ_i satisfy the property $\lambda_i \le \gamma$.*
- *$\mathrm{I} - \mathsf{TI}_\gamma(\mathcal{P})$ projects a quantum state into the subspace spanned by $\{|\psi_i\rangle\}$ whose eigenvalues λ_i satisfy the property $\lambda_i > \gamma$.*

The proof of the theorem above follows directly from Theorem 6.3 by considering the projective measurements $\mathsf{TI}_\gamma(\mathcal{P}) = \sum_{i:\lambda_i \le \gamma} E_{\lambda_i}$ and $\mathrm{I} - \mathsf{TI}_\gamma(\mathcal{P}) = \mathrm{I} - \sum_{i:\lambda_i > \gamma} E_{\lambda_i}$.

Finally, we also use the following *symmetric* variant of threshold implementations which were considered in [AKL+22, Theorem 2.6]. Here, the projective measurement determines whether the success probability is either close to $1/2$ or far from $1/2$.

Theorem 6.6 (Symmetric threshold implementation)**.** *Let $\gamma \in (0, 1/2)$ be a parameter and let $\mathcal{P} = (P,Q)$ be a two-outcome POVM, where P has an eigenbasis $\{|\psi_i\rangle\}$ with associated eigenvalues $\{\lambda_i\}$. Then, there exists a projective threshold implementation $(\mathsf{STI}_\gamma(\mathcal{P}), \mathrm{I} - \mathsf{STI}_\gamma(\mathcal{P}))$ such that*

- *$\mathsf{STI}_\gamma(\mathcal{P})$ projects a quantum state into the subspace spanned by $\{|\psi_i\rangle\}$ whose eigenvalues λ_i satisfy the property $|\lambda_i - \frac{1}{2}| \le \gamma$.*

- $\mathrm{I} - \mathsf{STI}_\gamma(\mathcal{P})$ *projects a quantum state into the subspace spanned by* $\{|\psi_i\rangle\}$ *whose eigenvalues* λ_i *satisfy the property* $|\lambda_i - \frac{1}{2}| > \gamma$.

The proof of the theorem above follows directly from Theorem 6.3 by considering the projective measurements $\mathsf{STI}_\gamma(\mathcal{P}) = \sum_{i:|\lambda_i - \frac{1}{2}| \leq \gamma} E_{\lambda_i}$ and $\mathrm{I} - \mathsf{STI}_\gamma(\mathcal{P}) = \mathrm{I} - \sum_{i:|\lambda_i - \frac{1}{2}| > \gamma} E_{\lambda_i}$.

Efficient measurements. The quantum measurements we described above can, in general, not be implemented efficiently. However, Zhandry [Zha20] showed that there exist so-called efficient *approximate implementations* which allow one to obtain approximate estimates of the success probability of a quantum program. In this section, we review some basic definitions and results that allow us to perform such measurmenents efficiently.

Definition 6.7 (Mixture of projective measurements)**.** *Let* $\mathcal{P} = \{\mathcal{P}_i\}_{i \in \mathcal{I}}$ *be a collection of binary outcome projective measurements* $\mathcal{P}_i = (P_i, Q_i)$ *over the same Hilbert space* $\mathcal{H}$*, and suppose that* P_i *corresponds to outcome* 1 *and* Q_i *corresponds to outcome* 0*. Let* D *be a distribution over the the index set* $\mathcal{I}$*. Then,* $\mathcal{P}_D = (P_D, Q_D)$ *is the following mixture of pojective measurements:*

$$P_D = \sum_{i \in \mathcal{I}} \Pr[i \leftarrow D]\, P_i \qquad \text{and} \qquad Q_D = \sum_{i \in \mathcal{I}} \Pr[i \leftarrow D]\, Q_i.$$

The following result is adapted from [Zha20, Theorem 6.2] and [ALL+21, Corollary 1].

Lemma 6.8 (Approximate threshold implementation)**.** *Let* $\mathcal{P}_D = (P_D, Q_D)$ *be a binary outcome POVM over Hilbert space* $\mathcal{H}$ *that is a mixture of projective measurements over some distribution* D*. Let* $\varepsilon, \delta, \gamma \in (0,1)$*. Then, there exists an efficient binary-outcome quantum algorithm* $\mathsf{ATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma}$*, interpreted as the POVM element corresponding to outcome* 1*, such that the following holds:*

- *For all quantum states* ρ*,* $\mathrm{Tr}[\mathsf{ATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma-\epsilon}\, \rho] \geq \mathrm{Tr}[\mathsf{TI}_\gamma(\mathcal{P}_D)\, \rho] - \delta$.
- *For all quantum states* ρ*, it holds that* $\mathrm{Tr}[\mathsf{TI}_{\gamma-2\varepsilon}(\mathcal{P}_D)\, \rho'] \geq 1 - 2\delta$*, where* ρ' *is the post-measurement state which results from applying the measurement* $\mathsf{ATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma}$ *to* ρ*.*
- *The expected running time to implement* $\mathsf{ATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma}$ *is proportional to* $\mathrm{poly}(1/\varepsilon, \log(1/\delta))$*, the time it takes to implement* P_D*, and the time it takes to sample from* D*.*

Finally, we use the following *symmetric* version of the approximate threshold implementation Lemma 6.8 which is a variant of [AKL+22, Theorem 2.8].

Lemma 6.9 (Symmetric approximate threshold implementation)**.** *Let* $\mathcal{P}_D = (P_D, Q_D)$ *be a binary outcome POVM over Hilbert space* $\mathcal{H}$ *that is a mixture of projective measurements over some distribution* D*. Let* $\gamma \in (0, 1/2)$ *and* $\varepsilon \in (0, \gamma/2)$*, and let* $\delta \in (0,1)$*. Let* D *be a distribution. Then, there exists an efficient binary-outcome quantum algorithm* $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma}$*, interpreted as the POVM element corresponding to outcome* 1*, such that the following holds:*

- *For all quantum states* ρ*,* $\mathrm{Tr}[\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma-\epsilon}\, \rho] \geq \mathrm{Tr}[\mathsf{STI}_\gamma(\mathcal{P}_D)\, \rho] - \delta$.
- *For all quantum states* ρ*, it holds that* $\mathrm{Tr}[\mathsf{STI}_{\gamma-2\varepsilon}(\mathcal{P}_D)\, \rho'] \geq 1 - 2\delta$*, where* ρ' *is the post-measurement state which results from applying the measurement* $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma}$ *to* ρ*.*
- *The expected running time to implement* $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma}$ *is proportional to* $\mathrm{poly}(1/\varepsilon, \log(1/\delta))$*, the time it takes to implement* P_D*, and the time it takes to sample from* D*.*

Useful Lemma. We are now ready to prove an important lemma. Roughly speaking, the lemma says the following. Suppose we have a bipartite state ρ on two registers R and AUX with the guarantee that conditioned on a binary outcome POVM succeeding on R, given the register AUX, a distinguisher can successfully distinguish two distributions $\mathcal{D}_0$ and $\mathcal{D}_1$. The lemma states that there is a distinguisher that can distinguish $\mathcal{D}_0$ and $\mathcal{D}_1$ *regardless of the outcome of the binary-outcome POVM on* R.

Lemma 6.10. *Let $\lambda \in \mathbb{N}$ be a parameter and let $\rho_{\mathrm{R},\mathrm{AUX}}$ be a quantum state on systems* R *and* AUX *of at most* $\mathrm{poly}(\lambda)$ *many qubits. Let D_0, D_1 be two efficiently samplable distributions with support $\mathcal{X}$. Let $\mathcal{D}$ be a* QPT *algorithm. Suppose that the following two properties hold:*

- *A (possibly inefficient) two-outcome POVM $\mathcal{M} = \{M_1, M_0\}$ succeeds on system* R *with probability at least*
$$\mathrm{Tr}[(M_1 \otimes I_{\mathrm{AUX}})\rho] \geq \frac{1}{p(\lambda)}$$
for some polynomial $p(\lambda)$.
- *the algorithm $\mathcal{D}$ succeeds at distinguishing D_0 from D_1 with advantage*
$$\left|\Pr\left[\mathcal{D}(x, \mathrm{AUX}) = b \,:\, \begin{smallmatrix} b \xleftarrow{\$} \{0,1\} \\ x \sim D_b \\ 1 \leftarrow \mathcal{M}(\mathrm{R}) \end{smallmatrix}\right] - \frac{1}{2}\right| \geq \frac{1}{q(\lambda)}.$$
for some polynomial $q(\lambda)$ conditioned on the measurement $\mathcal{M}$ succeeding on register R.

Then, there exists a QPT *algorithm $\tilde{\mathcal{D}}$ and a polynomial $\mu(\lambda)$ such that $\tilde{\mathcal{D}}$ succeeds at distinguishing D_0 and D_1 with advantage at least $1/\mu(\lambda)$ on the reduced system alone, i.e.*
$$\left|\Pr\left[\tilde{\mathcal{D}}(x, \mathrm{AUX}) = b \,:\, \begin{smallmatrix} b \xleftarrow{\$} \{0,1\} \\ x \sim D_b \end{smallmatrix}\right] - \frac{1}{2}\right| \geq \frac{1}{\mu(\lambda)},$$
where system AUX *corresponds to the reduced state $\rho_{\mathrm{AUX}} = \mathrm{Tr}_{\mathrm{R}}[\rho_{\mathrm{R},\mathrm{AUX}}]$.*

Proof. Consider the binary outcome POVM $\mathcal{P} = (P_{(D_0,D_1)}, Q_{(D_0,D_1)})$ with $Q_{(D_0,D_1)} = I - P_{(D_0,D_1)}$ which is the following mixture of projective measurents such that
$$P_{(D_0,D_1)} = \frac{\Pi_0 + \Pi_1}{2}$$
where Π_0, Π_1 are mixtures of two-outcome POVMs $\{\mathcal{P}_x\}$ that correspond to running $\mathcal{D}$ on samples x from D_0, D_1 and system AUX, and then measuring whether the output is 0 or 1, i.e.
$$\Pi_0 = \sum_{x \in \mathcal{X}} \Pr[x \leftarrow D_0]\, \mathcal{P}_x \qquad \text{and} \qquad \Pi_1 = \sum_{x \in \mathcal{X}} \Pr[x \leftarrow D_1]\, \mathcal{P}_x.$$
Let $\mathcal{P}$ have an eigenbasis $\{|\psi_i\rangle\}$ with eigenvalues $\{\lambda_i\}$. Without loss of generality we can assume that $\rho_{R,\mathrm{AUX}}$ is a pure state $|\psi\rangle = \sum_i \alpha_i |\psi\rangle$ in systems R and AUX. Moreover, we can write $|\psi\rangle$ as
$$|\psi\rangle = \sum_{i:\, |\lambda_i - \frac{1}{2}| \geq \frac{1}{q}} \alpha_i |\psi_i\rangle + \sum_{i:\, |\lambda_i - \frac{1}{2}| < \frac{1}{q}} \alpha_i |\psi_i\rangle .$$
Let $\varepsilon = 1/8p$, $\delta = 2^{-\lambda}$ and $\gamma = 1/2p$ be parameters. Consider the following distinguisher $\tilde{\mathcal{D}}$:

- Run the efficient approximate threshold implementation $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}$ from Lemma 6.9 on system Aux for the binary-outcome POVM given by $\mathcal{P}$.
- If the outcome of $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}$ is 1, then run $\mathcal{D}$ on the post-measurement system Aux, and output whatever $\mathcal{D}$ outputs. Otherwise, output a random bit.

Let us now analyze the success probability of $\tilde{\mathcal{D}}$. Because the two-outcome POVM $\mathcal{M}$ succeeds on system R with probability at least $\frac{1}{p}$ and because $\mathcal{D}$ succeeds with advantage at least $\frac{1}{q}$ on system Aux conditioned on $\mathcal{M}$ outputting 1, we have that $|\psi\rangle$ has weight at least $\frac{1}{p}$ on eigenvectors with eigenvalues λ_i such that $|\lambda_i - \frac{1}{2}| \geq \frac{1}{q}$. In other words,

$$\sum_{i:\,|\lambda_i-\frac{1}{2}|\geq\frac{1}{q}} |\alpha_i|^2 \geq \frac{1}{p}.$$

Therefore, the probability that $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}$ outputs 1 on system Aux is at least

$$\mathrm{Tr}\left[\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}(\mathrm{Aux})\right] \geq \frac{1}{p} - 2\delta = O\left(1/p\right).$$

Moreover, the post-measurement state in system $\tilde{\mathrm{Aux}}$ after getting outcome 1 has weight $1-2\delta$ on eigenvectors $\{|\psi_i\rangle\}$ such that $|\lambda_i - \frac{1}{2}| > \gamma - 2\epsilon$. Therefore, with probability at least $1-2\delta$, $\mathcal{D}$ has an advantage of at least $\gamma - 2\epsilon$ at outputting the correct bit when run on the collapsed post-measurement system $\tilde{\mathrm{Aux}}$.

However, if the measurement $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}$ on system Aux fails and outputs 0, then $\tilde{\mathcal{D}}$ succeeds with probability 1/2. Therefore, with overwhelming probability, $\tilde{\mathcal{D}}$ has advantage at least

$$\begin{aligned}
&\left|\Pr\left[\tilde{\mathcal{D}}(x,\mathrm{Aux}) = b \,:\, b \xleftarrow{\$} \{0,1\} \atop x\sim D_b\right] - \frac{1}{2}\right| \\
&= \left|\Pr\left[\mathcal{D}(x,\tilde{\mathrm{Aux}}) = b \,:\, b \xleftarrow{\$} \{0,1\} \atop x\sim D_b\right] \cdot \mathrm{Tr}\left[\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}(\mathrm{Aux})\right]\right. \\
&\qquad \left. + \frac{1}{2}\left(1 - \mathrm{Tr}\left[\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}(\mathrm{Aux})\right]\right) - \frac{1}{2}\right| \\
&= \mathrm{Tr}\left[\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}(\mathrm{Aux})\right] \cdot \left|\Pr\left[\mathcal{D}(x,\tilde{\mathrm{Aux}}) = b \,:\, b \xleftarrow{\$} \{0,1\} \atop x\sim D_b\right] - \frac{1}{2}\right| \\
&\geq (1/p - 2\delta)\cdot(\gamma - 2\varepsilon) \;\geq\; 1/\mathrm{poly}(\lambda).
\end{aligned}$$

Finally, we remark that the running time of the distinguisher $\tilde{\mathcal{D}}$ is proportional to the running time of $\mathcal{D}$ and $\mathrm{poly}(1/\varepsilon, \log(1/\delta))$, and hence it is efficient. □

6.3 Simultaneous Search-to-Decision Reduction with Quantum Auxiliary Input

Our first result concerns distinguishers with quantum auxiliary input that can distinguish between Dual-Regev samples and uniformly random samples with high probability. In Theorem 6.11, we give a search-to-decision reduction: we show that such distinguishers can be converted into a quantum extractor that can obtain a Dual-Regev secret key with overwhelming probability. We then state a strenghtening of this extraction property (which we call *Simultaneous Dual-Regev Extraction*)

in Theorem 6.16. Informally, this property states that extraction is possible even if additionally require that a *revocation* procedure succeeds on a separate register.

While we do not know how to prove Theorem 6.16 under standard assumptions, we prove that *Simultaneous Dual-Regev Extraction* holds assuming $\mathsf{LWE}/\mathsf{SIS}$ in the special case when revocation succeeds with overwhelming probability. This is captured by Theorem 6.19.

Search-to-decision reduction. We first show the following result.

Theorem 6.11 (Search-to-Decision Reduction with Quantum Auxiliary Input)**.** *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and let $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and let $\alpha, \beta \in (0,1)$ be noise ratios with $\beta/\alpha = 2^{o(n)}$ and $1/\alpha = 2^{o(n)} \cdot \sigma$. Let $\mathcal{A} = \{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}, \nu_\lambda)\}_{\lambda\in\mathbb{N}}$ be any non-uniform quantum algorithm consisting of a family of polynomial-sized quantum circuits*

$$\left\{\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}} : \mathcal{L}(\mathcal{H}_q^m \otimes \mathcal{H}_{B_\lambda}) \to \mathcal{L}(\mathcal{H}_{R_\lambda} \otimes \mathcal{H}_{\mathrm{AUX}_\lambda})\right\}_{\mathbf{A}\in\mathbb{Z}_q^{n\times m}, \mathbf{y}\in\mathbb{Z}_q^n}$$

and polynomial-sized advice states $\nu_\lambda \in \mathcal{D}(\mathcal{H}_{B_\lambda})$ which are independent of $\mathbf{A}$. Then, assuming the quantum hardness of the $\mathsf{LWE}^m_{n,q,\alpha q}$ assumption, the following holds for every QPT distinguisher $\mathcal{D}$. Suppose that there exists a function $\varepsilon(\lambda) = 1/\mathrm{poly}(\lambda)$ such that

$$\left|\Pr\left[1 \leftarrow \mathsf{SearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{SearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda, 1)\right]\right| = \varepsilon(\lambda).$$

$\underline{\mathsf{SearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}\left(1^\lambda, b\right)}$:

- If $b = 0$: output $\mathsf{lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 3.
- If $b = 1$: output $\mathsf{unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 4.

Figure 2: The experiment $\mathsf{SearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}\left(1^\lambda, b\right)$.

Then, there exists a quantum extractor $\mathcal{E}$ that takes as input $\mathbf{A}$, $\mathbf{y}$ and system AUX of the state $\rho_{\mathrm{R},\mathrm{AUX}}$ and outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\mathrm{poly}(\lambda, m, \sigma, q, 1/\varepsilon)$ such that

$$\Pr\left[\begin{matrix}\mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}})=\mathbf{x} \\ \wedge \\ \mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{\frac{m}{2}})\end{matrix} : \begin{matrix}\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m} \\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A},\sigma) \\ \rho_{\mathrm{R},\mathrm{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)\end{matrix}\right] \geq 1/\mathrm{poly}(\lambda).$$

Proof. Let $\lambda \in \mathbb{N}$ be the security parameter and let $\mathcal{A} = \{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}, \nu_\lambda)\}_{\mathbf{A}\in\mathbb{Z}_q^{n\times m}}$ be a non-uniform quantum algorithm. Suppose that $\mathcal{D}$ is a QPT distinguisher with advantage $\varepsilon = 1/\mathrm{poly}(\lambda)$.

To prove the claim, we consider the following sequence of hybrid distributions.

$\underline{\mathsf{lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)}$:

1. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Generate $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$.
3. Generate $\rho_{R,\,\mathrm{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)$.
4. Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n, \mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ and $e' \sim D_{\mathbb{Z},\beta q}$.
5. Generate $\rho_{\mathrm{R,AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)$.
6. Run $b' \leftarrow \mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal, \mathbf{s}^\intercal\mathbf{y} + e', \rho_{\mathrm{AUX}})$ on the reduced state. Output b'.

Figure 3: The distribution $\mathsf{lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)$.

H_0: This is the distribution $\mathsf{lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)$ in Figure 3.

H_1: This is the following distribution:

1. Sample a random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ and $e' \sim D_{\mathbb{Z},\beta q}$.
3. Sample a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}}$ and let $\mathbf{y} = \mathbf{A}\cdot\mathbf{x}_0 \pmod q$.
4. Run $\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{x}_0\rangle\langle\mathbf{x}_0| \otimes \nu_\lambda)$ to generate a state $\rho_{\mathrm{R,AUX}}$ in systems R and AUX.
5. Run the distinguisher $\mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal, \mathbf{s}^\intercal\mathbf{y} + e', \rho_{\mathrm{AUX}})$ on the reduced state ρ_{AUX}.

H_2 : This is the following distribution:

1. Sample a uniformly random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ and $e' \sim D_{\mathbb{Z},\beta q}$. Let $\mathbf{u} = \mathbf{A}^\intercal\mathbf{s} + \mathbf{e}$.
3. Sample a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}}$ and let $\mathbf{y} = \mathbf{A}\cdot\mathbf{x}_0 \pmod q$.
4. Run $\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{x}_0\rangle\langle\mathbf{x}_0| \otimes \nu_\lambda)$ to generate a state $\rho_{R,\mathrm{AUX}}$ in systems R and AUX.
5. Run the distinguisher $\mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 + e', \rho_{\mathrm{AUX}})$ on the reduced state ρ_{AUX}.

H_3 : This is the following distribution:

1. Sample a uniformly random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Sample $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and $e' \sim D_{\mathbb{Z},\beta q}$.
3. Sample a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}}$ and let $\mathbf{y} = \mathbf{A}\cdot\mathbf{x}_0 \pmod q$.
4. Run $\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{x}_0\rangle\langle\mathbf{x}_0| \otimes \nu_\lambda)$ to generate a state $\rho_{R,\mathrm{AUX}}$ in systems R and AUX.
5. Run the distinguisher $\mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 + e', \rho_{\mathrm{AUX}})$ on the reduced state ρ_{AUX}.

H_4: This is the following distribution:

1. Sample a uniformly random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Sample $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and $r \xleftarrow{\$} \mathbb{Z}_q$.
3. Sample a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and let $\mathbf{y} = \mathbf{A}\cdot\mathbf{x}_0 \pmod q$.
4. Run $\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{x}_0\rangle\langle\mathbf{x}_0| \otimes \nu_\lambda)$ to generate a state $\rho_{R,\text{AUX}}$ in systems R and AUX.
5. Run the distinguisher $\mathcal{D}(\mathbf{A},\mathbf{y},\mathbf{u},r,\rho_{\text{AUX}})$ on the reduced state ρ_{AUX}.

H_5: This is the distribution $\mathsf{unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ in Figure 4.

$\underline{\mathsf{unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)}$:

1. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Sample $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and $r \xleftarrow{\$} \mathbb{Z}_q$.
3. Run $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A},\sigma)$.
4. Generate $\rho_{R,\text{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)$.
5. Run $b' \leftarrow \mathcal{D}(\mathbf{A},\mathbf{y},\mathbf{u},r,\rho_{\text{AUX}})$ on the reduced state. Output b'.

Figure 4: The distribution $\mathsf{unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$.

We now show the following:

Claim 6.12. *Assuming* $\mathsf{LWE}^m_{n,q,\alpha q}$, *the hybrids* H_0 *and* H_1 *are computationally indistinguishable,*

$$\mathsf{H}_0 \approx_c \mathsf{H}_1.$$

Proof. Here, we invoke the *Gaussian-collapsing property* in Theorem 3.1 which states that the following samples are indistinguishable under $\mathsf{LWE}^m_{n,q,\alpha q}$,

$$\left(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m},\ |\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m\\ \mathbf{A}\mathbf{x}=\mathbf{y}}} \rho_\sigma(\mathbf{x})\ |\mathbf{x}\rangle\,,\ \mathbf{y}\in\mathbb{Z}_q^n\right) \approx_c \left(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m},\ |\mathbf{x}_0\rangle\,,\ \mathbf{A}\cdot\mathbf{x}_0 \in \mathbb{Z}_q^n\right)$$

where $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A},\sigma)$ and where $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ is a sample from the discrete Gaussian distribution. Because $\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}$ is a family efficient quantum algorithms, this implies that

$$\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda) \quad \approx_c \quad \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{x}_0\rangle\langle\mathbf{x}_0| \otimes \nu_\lambda),$$

for any polynomial-sized advice state $\nu_\lambda \in \mathcal{D}(\mathcal{H}_{B_\lambda})$ which is independent of $\mathbf{A}$. □

Claim 6.13. *Hybrids* H_1 *and* H_2 *are statistically indistinguishable. In other words,*

$$\mathsf{H}_1 \approx_s \mathsf{H}_2.$$

Proof. Here, we invoke the *noise flooding* property in Lemma 2.9 to argue that $\mathbf{e}^\intercal \mathbf{x}_0 \ll e'$ holds with overwhelming probability for our choice of parameters. Therefore, the distributions in H_1 and H_2 are computationally indistinguishable. □

Claim 6.14. *Assuming* $\mathsf{LWE}^m_{n,q,\alpha q}$*, the hybrids* H_2 *and* H_3 *are computationally indistinguishable,*

$$\mathsf{H}_2 \approx_c \mathsf{H}_3.$$

Proof. This follows from the $\mathsf{LWE}^m_{n,q,\alpha q}$ assumption since the reduction can sample $\mathbf{x}_0 \sim D_{\mathbb{Z}^m,\frac{\sigma}{\sqrt{2}}}$ itself and generate $\rho_{R,\mathrm{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{x}_0\rangle\langle\mathbf{x}_0| \otimes \nu_\lambda)$ on input $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and ν_λ. □

Finally, we show the following:

Claim 6.15. *Assuming* $\mathsf{LWE}^m_{n,q,\alpha q}$*, the hybrids* H_4 *and* H_5 *are computationally indistinguishable,*

$$\mathsf{H}_4 \approx_c \mathsf{H}_5.$$

Proof. Here, we invoke the *Gaussian-collapsing property* in Theorem 3.1 again. □

Recall that H_0 and H_5 can be distinguished with probability $\varepsilon = 1/\mathrm{poly}(\lambda)$. We proved that the hybrids H_0 and H_3 are computationally indistinguishable and moreover, hybrids H_4 and H_5 are computationally indistinguishable. As a consequence, it holds that hybrids H_3 and H_4 can be distinguished with probability at least $\varepsilon - \mathsf{negl}(\lambda)$.

We leverage this to obtain a Goldreich-Levin reduction. Consider the following distinguisher.

$\underline{\tilde{\mathcal{D}}\big(\mathbf{A},\mathbf{y},\mathbf{u},v,\rho\big)}$:

Input: $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$, $\mathbf{u} \in \mathbb{Z}_q^n$, $v \in \mathbb{Z}_q$ and $\rho \in L(\mathcal{H}_{\mathrm{AUX}})$.
Output: A bit $b' \in \{0,1\}$.

Procedure:

1. Sample $e' \sim D_{\mathbb{Z},\beta q}$.
2. Output $b' \leftarrow \mathcal{D}\big(\mathbf{A},\mathbf{y},\mathbf{u},v+e',\rho\big)$.

Figure 5: The distinguisher $\tilde{\mathcal{D}}\big(\mathbf{A},\mathbf{y},\mathbf{u},v,\rho\big)$.

Note that $r + e' \pmod q$ is uniform whenever $r \xleftarrow{\$} \mathbb{Z}_q$ and $e' \sim D_{\mathbb{Z},\beta q}$. Therefore, our previous argument shows that there exists a negligible function η such that:

$$\left|\Pr\left[\tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{u},\mathbf{u}^\intercal\mathbf{x}_0,\rho_{\mathrm{AUX}})=1 \,:\, \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m},\mathbf{u}\xleftarrow{\$}\mathbb{Z}_q^m\\ \mathbf{x}_0\sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}},\mathbf{y}\leftarrow\mathbf{A}\cdot\mathbf{x}_0 \pmod q\\ \rho_{R\,\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{x}_0\rangle\langle\mathbf{x}_0|\otimes\nu_\lambda)\end{array}\right]\right.$$
$$\left.-\Pr\left[\tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{u},\mathbf{r},\rho_{\mathrm{AUX}})=1 \,:\, \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ \mathbf{u}\xleftarrow{\$}\mathbb{Z}_q^m,r\xleftarrow{\$}\mathbb{Z}_q\\ \mathbf{x}_0\sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}},\mathbf{y}\leftarrow\mathbf{A}\cdot\mathbf{x}_0 \pmod q\\ \rho_{R\,\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{x}_0\rangle\langle\mathbf{x}_0|\otimes\nu_\lambda)\end{array}\right]\right| \geq \varepsilon - \eta(\lambda).$$

From Theorem 4.3, it follows that there exists a Goldreich-Levin extractor $\mathcal{E}$ running in time $T(\mathcal{E}) = \mathrm{poly}(\lambda, n, m, \sigma, q, 1/\varepsilon)$ that outputs a short vector in $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ with probability at least

$$\Pr\left[\begin{array}{c} \mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}})=\mathbf{x} \\ \bigwedge \\ \mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}}) \end{array} : \begin{array}{c} \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}, \mathbf{x}_0 \sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}} \\ \mathbf{y} \leftarrow \mathbf{A}\cdot\mathbf{x}_0 \pmod q \\ \rho_{\mathrm{R,AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{x}_0\rangle\langle\mathbf{x}_0| \otimes \nu_\lambda) \end{array}\right] \geq \mathrm{poly}(\varepsilon, 1/q).$$

Assuming the $\mathsf{LWE}^m_{n,q,\alpha q}$ assumption, we can invoke the Gaussian-collapsing property in Theorem 3.1 once again which implies that the quantum extractor $\mathcal{E}$ satisfies

$$\Pr\left[\begin{array}{c} \mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}})=\mathbf{x} \\ \bigwedge \\ \mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}}) \end{array} : \begin{array}{c} \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m} \\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A},\sigma) \\ \rho_{\mathrm{R,AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda) \end{array}\right] \geq \mathrm{poly}(\varepsilon, 1/q).$$

This proves the claim. □

Simultaneous sarch-to-decision reduction. Next, we give a strengthening of our result in Theorem 6.11 and state a *simultaneous* search-to-decision reduction with quantum auxiliary input which holds even if additionally require that a *revocation* procedure succeeds on a separate register.

To formalize the notion that revocation is applied on a separate register, we introduce a procedure called $\mathsf{IneffRevoke}$ which is defined in Figure 6. In Figure 7, we also introduce an inefficient variant of $\mathsf{QSampGauss}$ from Section 3.3 which does not require a trapdoor.

$\underline{\mathsf{IneffRevoke}(\mathbf{A}, \mathbf{y}, \sigma, \rho_{\mathrm{R}})}$:

Input: $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$ and $\rho \in L(\mathcal{H}_{\mathrm{R}})$.
Output: Accept ($\top$) or reject ($\bot$).

Procedure:

1. Apply the (inefficient) projective measurement

$$\left\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, I - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\right\}$$

where $|\psi_{\mathbf{y}}\rangle$ is the Gaussian coset state

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m: \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle\,.$$

2. If the measurement succeeds, output $\top$. Else, output $\bot$.

Figure 6: The procedure $\mathsf{IneffRevoke}(\mathbf{A}, \mathbf{y}, \sigma, \rho_{\mathrm{R}})$.

We use the following conjecture. We refer the reader to the introduction for an informal explanation of the conjecture below.

$\underline{\mathsf{IneffQSampGauss}(\mathbf{A}, \mathbf{y}, \sigma)}$:

Input: $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$ and $\sigma > 0$.
Output: Gaussian coset state $|\psi_{\mathbf{y}}\rangle$.

Procedure:

1. Prepare the Gaussian coset state

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m : \\ \mathbf{A}\mathbf{x} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x}) |\mathbf{x}\rangle .$$

2. Output $|\psi_{\mathbf{y}}\rangle$.

Figure 7: The procedure $\mathsf{IneffQSampGauss}(\mathbf{A}, \mathbf{y}, \sigma)$.

Conjecture 1. *Let $\lambda \in \mathbb{N}$. Then, there exist parameters (each parameterized by λ) such that $n \in \mathbb{N}$, q is a prime with $q = 2^{o(n)}$, $m \geq 2n \log q$, $\sqrt{8m} < \sigma < q/\sqrt{8m}$, $\alpha \in (0,1)$ with $1/\alpha = 2^{o(n)} \cdot \sigma$ for which the following holds: for any QPT $\mathcal{A}$ and QPT $\mathcal{C}$ (and for a fixed algorithm $\mathcal{B}$), and for any fixed $\mathrm{poly}(\lambda)$-sized quantum auxiliary input ν_λ (which depends on λ):*

$$\left|\Pr[1 \leftarrow \mathsf{Expt}_{\mathcal{A},\mathcal{B},\mathcal{C}}(1^\lambda, 0)] - \Pr[1 \leftarrow \mathsf{Expt}_{\mathcal{A},\mathcal{B},\mathcal{C}}(1^\lambda, 1)]\right| \leq \mathsf{negl}(\lambda),$$

where $\mathsf{Expt}_{\mathcal{A},\mathcal{B},\mathcal{C}}(1^\lambda, b)$ is the asymmetric cloning experiment in Figure 8.

Theorem 6.16 (Simultaneous Search-To-Decision Reduction with Quantum Auxiliary Input)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n \in \mathbb{N}$, q be a prime with $q = 2^{o(n)}$, $m \geq 2n \log q$, $\sqrt{8m} < \sigma < q/\sqrt{8m}$, and let $\alpha, \beta \in (0,1)$ with $\beta/\alpha = 2^{o(n)}$ with $1/\alpha = 2^{o(n)} \cdot \sigma$. Let $\mathcal{A} = \{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}, \nu_\lambda)\}_{\lambda \in \mathbb{N}}$ be any non-uniform quantum algorithm consisting of a family of polynomial-sized quantum circuits*

$$\left\{\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}} : \mathcal{L}(\mathcal{H}_q^m \otimes \mathcal{H}_{B_\lambda}) \to \mathcal{L}(\mathcal{H}_{R_\lambda} \otimes \mathcal{H}_{\mathrm{AUX}_\lambda})\right\}_{\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{y} \in \mathbb{Z}_q^n}$$

and polynomial-sized advice states $\nu_\lambda \in \mathcal{D}(\mathcal{H}_{B_\lambda})$ which are independent of $\mathbf{A}$. Then, assuming Conjecture 1 is true, the following holds for every QPT distinguisher $\mathcal{D}$. Suppose there exists a function $\varepsilon(\lambda) = 1/\mathrm{poly}(\lambda)$ such that

$$\Big|\Pr\left[1 \leftarrow \mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda, 1)\right]\Big| = \varepsilon(\lambda).$$

$\underline{\mathsf{Expt}_{\mathcal{A},\mathcal{B},\mathcal{C}}(1^\lambda, b)}$:

1. The challenger samples a random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and lets $\mathbf{y} = \mathbf{A}\cdot\mathbf{x}_0 \pmod q$. Then, the challenger runs $|\psi_{\mathbf{y}}\rangle \leftarrow \mathsf{IneffQSampGauss}(\mathbf{A}, \mathbf{y}, \sigma)$ and sends $(\mathbf{A}, \mathbf{y}, |\psi_{\mathbf{y}}\rangle)$ to $\mathcal{A}$.

2. $\mathcal{A}$ receives $(\mathbf{A}, \mathbf{y}, |\psi_{\mathbf{y}}\rangle)$ together with auxiliary input ν_λ and generates a state ρ_{BC} in systems BC, and sends B to $\mathcal{B}$ and C to $\mathcal{C}$.

3. The challenger sends $(\mathbf{A}, \mathbf{y}, \sigma)$ to $\mathcal{B}$ and, depending on the value of b, the challenger sends the following to $\mathcal{C}$:

   - if $b = 0$: the challenger samples $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$, lets $\mathbf{u} = \mathbf{A}^\intercal\mathbf{s} + \mathbf{e}$, and sends $(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 \pmod q)$ to $\mathcal{C}$.
   - if $b = 0$: the challenger samples a uniformly random vector $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and sends $(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 \pmod q)$ to $\mathcal{C}$.

4. $\mathcal{B}$ receives as input $(\mathbf{A}, \mathbf{y}, \sigma, \mathrm{B})$ and runs $\mathsf{IneffRevoke}(\mathbf{A}, \mathbf{y}, \sigma, \mathrm{B})$. Next, $\mathcal{B}$ outputs $\top$, if the measurement succeeds, else $\mathcal{B}$ outputs $\bot$.

5. $\mathcal{C}$ receives as input $(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 \pmod q, \mathrm{C})$ and outputs a bit b'.

6. The challenger outputs 1, if $\mathcal{B}$ outputs $\top$ and $\mathcal{C}$ outputs $b' = b$. This is also the outcome of the experiment.

Figure 8: The experiment for Conjecture 1.

$\underline{\mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}\left(1^\lambda, b\right)}$:

- If $b = 0$: output $\mathsf{simult.lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 10.
- If $b = 1$: output $\mathsf{simult.unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 11.

Figure 9: The experiment $\mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}\left(1^\lambda, b\right)$.

Then, there exists a quantum extractor $\mathcal{E}$ that takes as input $\mathbf{A}$, $\mathbf{y}$ and system Aux *of the state $\rho_{\mathrm{R},\mathrm{Aux}}$ and outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\mathrm{poly}(\lambda, m, \sigma, q, 1/\varepsilon)$ such that*

$$\Pr\left[\begin{array}{c}\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\mathrm{R})=\top \\ \bigwedge \\ \mathcal{E}(\mathbf{A},\mathbf{y},\mathrm{Aux}) \in \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array} : \begin{array}{c}\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m} \\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A},\sigma) \\ \rho_{\mathrm{R},\mathrm{Aux}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)\end{array}\right] \geq \mathrm{poly}(\varepsilon, 1/q).$$

$\underline{\mathsf{simult.lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)}$:

1. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Generate $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$.
3. Generate $\rho_{\mathrm{R},\,\mathrm{Aux}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)$.
4. Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n, \mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ and $e' \sim D_{\mathbb{Z},\beta q}$.
5. Generate $\rho_{\mathrm{R},\mathrm{Aux}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)$.
6. Run $\mathsf{IneffRevoke}(\mathbf{A}, \mathbf{y}, \sigma, \cdot)$ on system R. If it outputs $\top$, continue. Otherwise, output $\mathsf{Invalid}$.
7. Run $b' \leftarrow \mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal, \mathbf{s}^\intercal\mathbf{y} + e', \cdot)$ on system Aux. Output b'.

Figure 10: The distribution $\mathsf{simult.lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$.

Proof. Let $\lambda \in \mathbb{N}$ be the security parameter and let $\mathcal{A} = \{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}, \nu_\lambda)\}_{\mathbf{A}\in\mathbb{Z}_q^{n\times m}}$ be a non-uniform quantum algorithm. Suppose that $\mathcal{D}$ is a QPT distinguisher with advantage $\varepsilon = 1/\mathrm{poly}(\lambda)$.

To prove the claim, we consider the following sequence of hybrid distributions.

H_0: This is the distribution $\mathsf{simult.lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ in Figure 10.

$\underline{\mathsf{simult.unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)}$:

1. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Generate $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$.
3. Sample $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and $r \xleftarrow{\$} \mathbb{Z}_q$.
4. Generate $\rho_{R,\text{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)$.
5. Run $\mathsf{IneffRevoke}(\mathbf{A}, \mathbf{y}, \sigma, \cdot)$ on system R. If it outputs $\top$, continue. Otherwise, output $\mathsf{Invalid}$.
6. Run $b' \leftarrow \mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{u}, r, \cdot)$ on system AUX. Output b'.

Figure 11: The distribution $\mathsf{simult.unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$.

H_1: This is the following distribution:

1. Sample a random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Sample a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and let $\mathbf{y} = \mathbf{A}\cdot\mathbf{x}_0 \pmod q$.
3. Run $|\psi_{\mathbf{y}}\rangle \leftarrow \mathsf{IneffQSampGauss}(\mathbf{A}, \mathbf{y}, \sigma)$.
4. Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ and $e' \sim D_{\mathbb{Z},\beta q}$. Let $\mathbf{u} = \mathbf{A}^\intercal\mathbf{s} + \mathbf{e}$.
5. Run $\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)$ to generate a state $\rho_{\text{R,AUX}}$ in systems R and AUX.
6. Run $\mathsf{IneffRevoke}(\mathbf{A}, \mathbf{y}, \sigma, \cdot)$ on R. If it outputs $\top$, continue. Otherwise, output $\mathsf{Invalid}$.
7. Run the distinguisher $\mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 + e', \cdot)$ on system AUX. Output b'.

H_2 : This is the following distribution:

1. Sample a random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Sample a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and let $\mathbf{y} = \mathbf{A}\cdot\mathbf{x}_0 \pmod q$.
3. Run $|\psi_{\mathbf{y}}\rangle \leftarrow \mathsf{IneffQSampGauss}(\mathbf{A}, \mathbf{y}, \sigma)$.
4. Sample $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and $e' \sim D_{\mathbb{Z},\beta q}$.
5. Run $\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_\lambda)$ to generate a state $\rho_{\text{R,AUX}}$ in systems R and AUX.
6. Run $\mathsf{IneffRevoke}(\mathbf{A}, \mathbf{y}, \sigma, \cdot)$ on R. If it outputs $\top$, continue. Otherwise, output $\mathsf{Invalid}$.
7. Run the distinguisher $\mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 + e', \cdot)$ on system AUX. Output b'.

H_3: This is the distribution $\mathsf{simult.unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 11.

We now show the following:

Claim 6.17. *Hybrids* H_0 *and* H_1 *are statistically indistinguishable. In other words,*

$$\mathsf{H}_0 \approx_s \mathsf{H}_1.$$

Proof. Here, we invoke the *noise flooding* property in Lemma 2.9 to argue that $\mathbf{e}^\intercal \mathbf{x}_0 \ll e'$ holds with overwhelming probability for our choice of parameters. Therefore, the distributions in H_0 and H_1 are computationally indistinguishable. □

Claim 6.18. *Assuming Conjecture 1 holds for our choice of parameters, the hybrids* H_1 *and* H_2 *are computationally indistinguishable,*

$$\mathsf{H}_1 \approx_c \mathsf{H}_2.$$

Proof. This follows directly from Conjecture 1 which allows us to invoke a variant of LWE assumption, even if the procedure $\mathsf{IneffRevoke}$ is applied on a separate register. Here, we rely on the fact that during the reduction, we can simply sample $e' \sim D_{\mathbb{Z},\beta q}$ to produce an identically distributed challenge distribution. □

Recall that H_0 and H_3 can be distinguished with probability $\varepsilon = 1/\mathrm{poly}(\lambda)$. We proved that the hybrids H_0 and H_2 are computationally indistinguishable. As a consequence, it holds that hybrids H_2 and H_3 can be distinguished with probability at least $\varepsilon - \mathsf{negl}(\lambda)$.

We leverage this to obtain a Goldreich-Levin reduction. Consider the following distinguisher.

$\underline{\tilde{\mathcal{D}}\big(\mathbf{A},\mathbf{y},\mathbf{u},v,\rho\big)}$:

Input: $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$, $\mathbf{u} \in \mathbb{Z}_q^n$, $v \in \mathbb{Z}_q$ and $\rho \in L(\mathcal{H}_{\mathrm{AUX}})$.
Output: A bit $b' \in \{0,1\}$.

Procedure:

1. Sample $e' \sim D_{\mathbb{Z},\beta q}$.
2. Output $b' \leftarrow \mathcal{D}\big(\mathbf{A},\mathbf{y},\mathbf{u},v+e',\rho\big)$.

Figure 12: The distinguisher $\tilde{\mathcal{D}}\big(\mathbf{A},\mathbf{y},\mathbf{u},v,\rho\big)$.

Note that $r + e' \pmod q$ is uniform whenever $r \xleftarrow{\$} \mathbb{Z}_q$ and $e' \sim D_{\mathbb{Z},\beta q}$. Therefore, our previous argument shows that there exists a negligible function η such that:

$$\left|\Pr\left[\begin{array}{c}\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\mathrm{R})=\top \\ \wedge \\ \tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{u},\mathbf{u}^\intercal\mathbf{x}_0,\mathrm{AUX})=1\end{array} : \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m},\mathbf{u}\xleftarrow{\$}\mathbb{Z}_q^m \\ \mathbf{x}_0\sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}},\mathbf{y}=\mathbf{A}\mathbf{x}_0 \pmod q \\ |\psi_{\mathbf{y}}\rangle\leftarrow\mathsf{IneffQSampGauss}(\mathbf{A},\mathbf{y},\sigma) \\ \rho_{\mathrm{R},\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right] - \right.$$

$$\left.\Pr\left[\begin{array}{c}\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\mathrm{R})=\top \\ \wedge \\ \tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{u},r,\mathrm{AUX})=1\end{array} : \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m} \\ \mathbf{u}\xleftarrow{\$}\mathbb{Z}_q^m,r\xleftarrow{\$}\mathbb{Z}_q \\ \mathbf{x}_0\sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}},\mathbf{y}=\mathbf{A}\mathbf{x}_0 \pmod q \\ |\psi_{\mathbf{y}}\rangle\leftarrow\mathsf{IneffQSampGauss}(\mathbf{A},\mathbf{y},\sigma) \\ \rho_{\mathrm{R},\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]\right| \geq \varepsilon - \eta(\lambda).$$

From Theorem 4.3, it follows that there exists a Goldreich-Levin extractor $\mathcal{E}$ running in time $T(\mathcal{E}) = \mathrm{poly}(\lambda, n, m, \sigma, q, 1/\varepsilon)$ that outputs a short vector in $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ with probability at least

$$\Pr\left[\begin{array}{c}\mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}})=\mathbf{x}\\ \wedge\\ \mathbf{x}\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array} : \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ \mathbf{x}_0\sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}},\mathbf{y}=\mathbf{A}\mathbf{x}_0 \pmod q\\ |\psi_{\mathbf{y}}\rangle\leftarrow\mathsf{IneffQSampGauss}(\mathbf{A},\mathbf{y},\sigma)\\ \rho_{\mathrm{R,AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]\geq\mathrm{poly}(\varepsilon,1/q).$$

This proves the claim. □

Finally, we give a proof of the *simultaneous* search-to-decision reduction (Theorem 6.16) from standard assumptions without Conjecture 1 in the special case when revocation succeeds with overwhelming probability.

Theorem 6.19. *Let $n\in\mathbb{N}$. Let q be a prime with $q=2^{o(n)}$, $m\geq 2n\log q$, $\sqrt{8m}<\sigma<q/\sqrt{8m}$, and let $\alpha,\beta\in(0,1)$ with $\beta/\alpha=2^{o(n)}$ with $1/\alpha=2^{o(n)}\cdot\sigma$. Let $\mathcal{A}=\{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}},\nu_\lambda)\}_{\lambda\in\mathbb{N}}$ be any non-uniform quantum algorithm consisting of a family of polynomial-sized quantum circuits*

$$\left\{\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}:\mathcal{L}(\mathcal{H}_q^m\otimes\mathcal{H}_{B_\lambda})\to\mathcal{L}(\mathcal{H}_{R_\lambda}\otimes\mathcal{H}_{\mathrm{AUX}_\lambda})\right\}_{\mathbf{A}\in\mathbb{Z}_q^{n\times m},\mathbf{y}\in\mathbb{Z}_q^n}$$

and polynomial-sized advice states $\nu_\lambda\in\mathcal{D}(\mathcal{H}_{B_\lambda})$ which are independent of $\mathbf{A}$ such that

$$\Pr\left[\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\rho_{\mathrm{R}})=\top : \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\mathrm{R,AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]=1-\nu(\lambda),$$

for some negligle function $\nu(\lambda)$. Then, assuming the quantum hardness of the $\mathsf{LWE}^m_{n,q,\alpha q}$ assumption, the following holds for every QPT distinguisher $\mathcal{D}$. Suppose that there exists a function $\varepsilon(\lambda)=1/\mathrm{poly}(\lambda)$ such that

$$\Big|\Pr\left[1\leftarrow\mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda,0)\right]-\Pr\left[1\leftarrow\mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda,1)\right]\Big|=\varepsilon(\lambda).$$

Then, there exists a quantum extractor $\mathcal{E}$ that takes as input $\mathbf{A}$, $\mathbf{y}$ and system AUX of the state $\rho_{\mathrm{R,AUX}}$ and outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\mathrm{poly}(\lambda,m,\sigma,q,1/\varepsilon)$ such that

$$\Pr\left[\begin{array}{c}\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\mathrm{R})=\top\\ \wedge\\ \mathcal{E}(\mathbf{A},\mathbf{y},\mathrm{AUX})\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array} : \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\mathrm{R,AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]\geq\mathrm{poly}(\varepsilon,1/q).$$

Proof. By assumption, there exists an adversary $(\mathcal{A},\mathcal{D})$ such that $\mathsf{Adv}(\mathcal{A},\mathcal{D})=\varepsilon(\lambda)$, where

$$\mathsf{Adv}(\mathcal{A},\mathcal{D})=\Bigg|\Pr\left[\begin{array}{c}\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\mathrm{R})=\top\\ \wedge\\ \mathcal{D}(\mathbf{A},\mathbf{y},\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal,\mathbf{s}^\intercal\mathbf{y}+e',\mathrm{AUX})=1\end{array} : \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m},\mathbf{s}\xleftarrow{\$}\mathbb{Z}_q^n\\ \mathbf{e}\sim D_{\mathbb{Z}^m,\alpha q},e'\sim D_{\mathbb{Z},\beta q}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\mathrm{R,AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]-\Pr\left[\begin{array}{c}\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\mathrm{R})=\top\\ \wedge\\ \mathcal{D}(\mathbf{A},\mathbf{y},\mathbf{u},r,\mathrm{AUX})=1\end{array} : \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ \mathbf{u}\xleftarrow{\$}\mathbb{Z}_q^m,r\xleftarrow{\$}\mathbb{Z}_q\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\mathrm{R,AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]\Bigg|.$$

We can now invoke Lemma 6.10 to argue that there exists a QPT distinguisher $\tilde{\mathcal{D}}$ (that internally runs $\mathcal{D}$) and succeeds on the reduced system Aux alone, i.e.

$$\left|\Pr\left[\tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal,\mathbf{s}^\intercal\mathbf{y}+e',\rho_{\text{Aux}})=1 \;:\; \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m},\mathbf{s}\xleftarrow{\$}\mathbb{Z}_q^n\\ \mathbf{e}\sim D_{\mathbb{Z}^m,\alpha q},e'\sim D_{\mathbb{Z},\beta q}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\text{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]-\right.$$
$$\left.\Pr\left[\tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{u},r,\rho_{\text{Aux}})=1 \;:\; \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ \mathbf{u}\xleftarrow{\$}\mathbb{Z}_q^m,r\xleftarrow{\$}\mathbb{Z}_q\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\text{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]\right|=\bar{\varepsilon}(\lambda),$$

for some $\bar{\varepsilon}=1/\text{poly}(\lambda)$. In other words, the QPT algorithm $\tilde{\mathcal{D}}$ can successfully predict whether it has received a Dual-Regev sample or a uniformly random sample. Therefore, we can now invoke Theorem 6.11 to argue there exists an extractor $\mathcal{E}$ that takes as input $\mathbf{A}$, $\mathbf{y}$ and system Aux of the state $\rho_{\text{R,Aux}}$ and outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\text{poly}(\lambda,m,\sigma,q,1/\bar{\varepsilon})$ such that

$$\Pr\left[\begin{array}{c}\mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\text{Aux}})=\mathbf{x}\\ \wedge\\ \mathbf{x}\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array}\;:\; \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\text{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]\geq\text{poly}(\bar{\varepsilon},1/q).$$

Recall also that, by assumption, revocation succeeds with overwhelming probability, i.e.,

$$\Pr\left[\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\rho_{\text{R}})=\top \;:\; \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\text{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]=1-\mathsf{negl}(\lambda).$$

Using Bonferroni's inequality, we can argue that

$$\begin{aligned}
&\Pr\left[\begin{array}{c}\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\text{R})=\top\\ \wedge\\ \mathcal{E}(\mathbf{A},\mathbf{y},\text{Aux})\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array}\;:\; \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\text{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]\\
&\geq\Pr\left[\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\rho_{\text{R}})=\top \;:\; \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\text{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]\\
&\quad+\Pr\left[\mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\text{Aux}})\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{m/2}) \;:\; \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\text{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]-1\\
&\geq\Pr\left[\mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\text{Aux}})\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{m/2}) \;:\; \begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\text{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\nu_\lambda)\end{array}\right]-\mathsf{negl}(\lambda)\\
&\geq\text{poly}(\varepsilon,1/q).
\end{aligned}$$

This proves the claim.

□

6.4 Distinct Pair Extraction

The following lemma allows us to analyze the probability of simultaneously extracting two distinct preimages in terms of the success probability of revocation and the success probability of extracting a preimage from the adversary's state.

Lemma 6.20 (Distinct pair extraction)**.** *Let $\rho \in \mathcal{D}(\mathcal{H}_X \otimes \mathcal{H}_Y)$ be an any density matrix, for some Hilbert spaces $\mathcal{H}_X$ and $\mathcal{H}_Y$. Let $|\psi\rangle = \sum_{x\in\mathcal{S}} \alpha_x |x\rangle \in \mathcal{H}_X$ be any state supported on a subset $\mathcal{S} \subseteq \mathcal{X}$, and let $\mathbf{\Pi} = |\psi\rangle\langle\psi|$ denote its associated projection. Let $\mathbf{\Pi}_\mathcal{S}$ be the projector onto $\mathcal{S}$ with*

$$\mathbf{\Pi}_\mathcal{S} = \sum_{x\in\mathcal{S}} |x\rangle\langle x|.$$

Let $\mathcal{E} : \mathcal{L}(\mathcal{H}_Y) \to \mathcal{L}(\mathcal{H}_{X'})$ be any CPTP *map of the form*

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

for some isometry $V_{Y\to X'E}$. Consider the measurement specified by

$$\mathbf{\Gamma} = \sum_{x,x'\in\mathcal{S}:x\neq x'} |x\rangle\langle x|_X \otimes V_{Y\to X'E}^\dagger (|x'\rangle\langle x'|_{X'} \otimes I_E) V_{Y\to X'E}.$$

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

$$\mathrm{Tr}[\mathbf{\Gamma}\rho] \;\geq\; \left(1 - \max_{x\in\mathcal{S}} |\alpha_x|^2\right) \cdot \mathrm{Tr}[\mathbf{\Pi}\rho_X] \cdot \mathrm{Tr}\left[\mathbf{\Pi}_\mathcal{S}\, \mathcal{E}_{Y\to X'}(\sigma)\right],$$

where $\sigma = \mathrm{Tr}[(\mathbf{\Pi}\otimes I)\rho]^{-1} \cdot \mathrm{Tr}_X[(\mathbf{\Pi}\otimes I)\rho]$ is a reduced state in system Y.

Proof. Because the order in which we apply $\mathbf{\Gamma}$ and $(\mathbf{\Pi}\otimes I)$ does not matter, we have the inequality

$$\mathrm{Tr}\left[\mathbf{\Gamma}\rho\right] \geq \mathrm{Tr}\left[(\mathbf{\Pi}\otimes I)\,\mathbf{\Gamma}\rho\right] = \mathrm{Tr}\left[(\mathbf{\Pi}\otimes I)\,\mathbf{\Gamma}\rho\,(\mathbf{\Pi}\otimes I)\right] = \mathrm{Tr}\left[\mathbf{\Gamma}(\mathbf{\Pi}\otimes I)\rho\,(\mathbf{\Pi}\otimes I)\right]. \tag{1}$$

Notice also that $(\mathbf{\Pi}\otimes I)\rho(\mathbf{\Pi}\otimes I)$ lies in the image of $(\mathbf{\Pi}\otimes I)$ with $\mathbf{\Pi} = |\psi\rangle\langle\psi|$, and thus

$$(\mathbf{\Pi}\otimes I)\rho(\mathbf{\Pi}\otimes I) = \mathrm{Tr}[(\mathbf{\Pi}\otimes I)\rho] \cdot (|\psi\rangle\langle\psi| \otimes \sigma), \tag{2}$$

for some $\sigma \in \mathcal{D}(\mathcal{H}_Y)$. Putting everything together, we get that

$$\begin{aligned}
\mathrm{Tr}\left[\mathbf{\Gamma}\rho\right] &\geq \mathrm{Tr}\left[\mathbf{\Gamma}(\mathbf{\Pi}\otimes I)\rho\,(\mathbf{\Pi}\otimes I)\right] && \text{(using inequality (1))}\\
&= \mathrm{Tr}[(\mathbf{\Pi}\otimes I)\rho] \cdot \mathrm{Tr}\left[\mathbf{\Gamma}\left(|\psi\rangle\langle\psi| \otimes \sigma\right)\right] && \text{(using equation (2))}\\
&= \mathrm{Tr}[\mathbf{\Pi}\rho_X] \cdot \mathrm{Tr}\left[\sum_{x,x'\in\mathcal{S}:x\neq x'} |x\rangle\langle x|_X \otimes V_{Y\to X'E}^\dagger \left(|x'\rangle\langle x'|_{X'} \otimes I_E\right) V_{Y\to X'E} \left(|\psi\rangle\langle\psi| \otimes \sigma\right)\right]\\
&= \mathrm{Tr}[\mathbf{\Pi}\rho_X] \cdot \sum_{x'\in\mathcal{S}} \left(\sum_{x\in\mathcal{S}:x\neq x'} |\langle x|\psi\rangle|^2\right) \mathrm{Tr}\left[V_{Y\to X'E}^\dagger (|x'\rangle\langle x'|_{X'} \otimes I_E) V_{Y\to X'E}\, \sigma\right]\\
&= \mathrm{Tr}[\mathbf{\Pi}\rho_X] \cdot \sum_{x'\in\mathcal{S}} \left(1 - |\alpha_{x'}|^2\right) \mathrm{Tr}\left[(|x'\rangle\langle x'|_{X'} \otimes I_E) V_{Y\to X'E}\, \sigma\, V_{Y\to X'E}^\dagger\right]\\
&\geq \mathrm{Tr}[\mathbf{\Pi}\rho_X] \cdot \left(1 - \max_{x\in\mathcal{S}} |\alpha_x|^2\right) \cdot \sum_{x'\in\mathcal{S}} \mathrm{Tr}\left[(|x'\rangle\langle x'|_{X'} \otimes I_E) V_{Y\to X'E}\, \sigma\, V_{Y\to X'E}^\dagger\right]\\
&= \mathrm{Tr}[\mathbf{\Pi}\rho_X] \cdot \left(1 - \max_{x\in\mathcal{S}} |\alpha_x|^2\right) \cdot \sum_{x'\in\mathcal{S}} \mathrm{Tr}\left[|x'\rangle\langle x'|_{X'}\, \mathrm{Tr}_E\left[V_{Y\to X'E}\, \sigma\, V_{Y\to X'E}^\dagger\right]\right]\\
&= \mathrm{Tr}[\mathbf{\Pi}\rho_X] \cdot \left(1 - \max_{x\in\mathcal{S}} |\alpha_x|^2\right) \cdot \mathrm{Tr}\left[\mathbf{\Pi}_\mathcal{S}\, \mathcal{E}_{Y\to X'}(\sigma)\right].
\end{aligned}$$

This proves the claim. □

6.5 Proof of Theorem 6.1

Proof. Let $\mathcal{A}$ be a QPT adversary and suppose that

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}_{\mathcal{A}}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}_{\mathcal{A}}(1^\lambda, 1)\right]\right| = \varepsilon(\lambda),$$

for some $\varepsilon(\lambda)$ with respect to $\mathsf{Expt}_{\mathcal{A}}(1^\lambda, b)$ in Figure 13. We show that $\varepsilon(\lambda)$ is negligible.

Suppose for the sake of contradiction that $\epsilon(\lambda)$ is non-negligible. Using the equivalence between prediction advantage and distinguishing advantage, we can write

$$2 \cdot \left|\Pr\left[b \leftarrow \mathsf{Expt}_{\mathcal{A}}(1^\lambda, b) \ : \ b \xleftarrow{\$} \{0,1\}\right] - \frac{1}{2}\right| = \varepsilon(\lambda).$$

We show that we can use $\mathcal{A}$ to break the $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ problem. Without loss of generality, we assume that $\mathcal{A}$ submits the plaintext $x = 0$. By the assumption that revocation succeeds with overwhelming probability and since $\epsilon(\lambda) \geq 1/\mathrm{poly}(\lambda)$, we can use Theorem 6.19 to argue that there exists a quantum Goldreich-Levin extractor $\mathcal{E}$ that takes as input $\mathbf{A}$, $\mathbf{y}$ and system Aux of the state $\rho_{R,\mathrm{Aux}}$ and outputs a short vector in the coset $\Lambda^{\mathbf{y}}_q(\mathbf{A})$ in time $\mathrm{poly}(\lambda, m, \sigma, q, 1/\varepsilon)$ such that

$$\Pr\left[\begin{array}{c}\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\mathrm{R})=\top \\ \wedge \\ \mathcal{E}(\mathbf{A},\mathbf{y},\mathrm{Aux}) \in \Lambda^{\mathbf{y}}_q(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array} : \begin{array}{c}\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m} \\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A},\sigma) \\ \rho_{R,\mathrm{Aux}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|)\end{array}\right] \geq \mathrm{poly}(\varepsilon, 1/q).$$

Here, we rely on the correctness of GenTrap in Theorem 2.13 and QSampGauss in Theorem 3.3. Consider the following procedure in Algorithm 3.

Algorithm 3: SIS_Solver($\mathbf{A}$)

Input: Matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$.
Output: Vector $\mathbf{x} \in \mathbb{Z}^m$.

1 Generate a Gaussian state $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ with

$$|\psi_{\mathbf{y}}\rangle \ = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x}) \ |\mathbf{x}\rangle$$

for some vector $\mathbf{y} \in \mathbb{Z}_q^n$.

2 Run $\mathcal{A}$ to generate a bipartite state $\rho_{R\,\mathrm{Aux}}$ in systems $\mathcal{H}_{\mathrm{R}} \otimes \mathcal{H}_{\mathrm{Aux}}$ with $\mathcal{H}_{\mathrm{R}} = \mathcal{H}_q^m$.

3 Measure system R in the computational basis, and let $\mathbf{x}_0 \in \mathbb{Z}_q^n$ denote the outcome.

4 Run the quantum Goldreich-Levin extractor $\mathcal{E}(\mathbf{A}, \mathbf{y}, \rho_{\mathrm{Aux}})$ from Theorem 6.16, where ρ_{Aux} is the reduced state in system $\mathcal{H}_{\mathrm{Aux}}$, and let $\mathbf{x}_1 \in \mathbb{Z}_q^n$ denote the outcome.

5 Output the vector $\mathbf{x} = \mathbf{x}_1 - \mathbf{x}_0$.

To conclude the proof, we show that SIS_Solver($\mathbf{A}$) in Algorithm 3 breaks the $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ problem whenever $\varepsilon(\lambda) = 1/\mathrm{poly}(\lambda)$. In order to guarantee that SIS_Solver($\mathbf{A}$) is successful, we use the distinct pair extraction result of Lemma 6.20. This allows us to analyze the probability of simultaneously extracting two distinct short pre-images $\mathbf{x}_0 \neq \mathbf{x}_1$ such that $\mathbf{A}\mathbf{x}_0 = \mathbf{y} = \mathbf{A}\mathbf{x}_1 \pmod q$ – both in terms of the success probability of revocation and the success probability of extracting a

$\underline{\mathsf{Expt}_{\mathcal{A}}(1^\lambda, b)}$:

1. The challenger samples $(\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ and generates

$$|\psi_{\mathbf{y}}\rangle \;=\; \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x})\;|\mathbf{x}\rangle\,,$$

   for some $\mathbf{y} \in \mathbb{Z}_q^n$, by running $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$. The challenger lets $\mathsf{MSK} \leftarrow \mathsf{td}_{\mathbf{A}}$ and $\mathsf{PK} \leftarrow (\mathbf{A}, \mathbf{y})$ and sends $\rho_{\mathsf{SK}} \leftarrow |\psi_{\mathbf{y}}\rangle$ to the adversary $\mathcal{A}$.

2. $\mathcal{A}$ generates a (possibly entangled) bipartite state $\rho_{R,\mathrm{AUX}}$ in systems $\mathcal{H}_R \otimes \mathcal{H}_{\mathrm{AUX}}$ with $\mathcal{H}_R = \mathcal{H}_q^m$, returns system R and holds onto the auxiliary system AUX.

3. The challenger runs $\mathsf{Revoke}(\mathsf{PK}, \mathsf{MSK}, \rho_R)$, where ρ_R is the reduced state in system R. If the outcome is $\top$, the game continues. Otherwise, output $\mathsf{Invalid}$.

4. $\mathcal{A}$ submits a plaintext bit $\mu \in \{0, 1\}$.

5. The challenger does the following depending on $b \in \{0, 1\}$:

   - if $b = 0$: the challenger samples a vector $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ and errors $\mathbf{e} \sim D_{\mathbb{Z}^m,\,\alpha q}$ and $e' \sim D_{\mathbb{Z},\,\beta q}$, and sends a Dual-Regev encryption of $\mu \in \{0, 1\}$ to $\mathcal{A}$:

$$\mathsf{CT} = \left(\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal, \mathbf{s}^\intercal\mathbf{y} + e' + \mu \cdot \lfloor\frac{q}{2}\rfloor\right) \in \mathbb{Z}_q^m \times \mathbb{Z}_q.$$

   - if $b = 1$: the challenger samples $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and $r \xleftarrow{\$} \mathbb{Z}_q$ uniformly at random and sends the following pair to $\mathcal{A}$:

$$(\mathbf{u}, r) \in \mathbb{Z}_q^m \times \mathbb{Z}_q.$$

6. $\mathcal{A}$ returns a bit $b' \in \{0, 1\}$.

Figure 13: The key-revocable security experiment according to Definition 5.3.

pre-image from the adversary's state ρ_{AUX} in system $\mathcal{H}_{\text{AUX}}$. Assuming that $\mathbf{x}_0, \mathbf{x}_1$ are distinct short pre-images such that $\|\mathbf{x}_0\| \leq \sigma\sqrt{\frac{m}{2}}$ and $\|\mathbf{x}_1\| \leq \sigma\sqrt{\frac{m}{2}}$, it then follows that the vector $\mathbf{x} = \mathbf{x}_1 - \mathbf{x}_0$ output by $\mathsf{SIS\_Solver}(\mathbf{A})$ has norm at most $\sigma\sqrt{2m}$, and thus yields a solution to $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$.

We remark that the state $|\psi_{\mathbf{y}}\rangle$ prepared by Algorithm 3 is not normalized for ease of notation. Note that the tail bound in Lemma 2.6 implies that (the normalized variant of) $|\psi_{\mathbf{y}}\rangle$ is within negligible trace distance of the state with support $\{\mathbf{x} \in \mathbb{Z}_q^m : \|\mathbf{x}\| \leq \sigma\sqrt{\frac{m}{2}}\}$. Therefore, for the sake of Lemma 6.20, we can assume that $|\psi_{\mathbf{y}}\rangle$ is a normalized state of the form

$$|\psi_{\mathbf{y}}\rangle = \left(\sum_{\substack{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \leq \sigma\sqrt{\frac{m}{2}} \\ \mathbf{Az} = \mathbf{y} \pmod q}} \rho_{\frac{\sigma}{\sqrt{2}}}(\mathbf{z}) \right)^{-\frac{1}{2}} \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m, \|\mathbf{x}\| \leq \sigma\sqrt{\frac{m}{2}} \\ \mathbf{Ax} = \mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \, .$$

Before we analyze Algorithm 3, we first make two technical remarks. First, since $\sigma \geq \omega(\sqrt{\log m})$, it follows from Lemma 2.10 that, for any full-rank $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and for any $\mathbf{y} \in \mathbb{Z}_q^n$, we have

$$\max_{\substack{\mathbf{x} \in \mathbb{Z}_q^m, \|\mathbf{x}\| \leq \sigma\sqrt{\frac{m}{2}} \\ \mathbf{Ax} = \mathbf{y} \pmod q}} \left\{ \frac{\rho_{\frac{\sigma}{\sqrt{2}}}(\mathbf{x})}{\sum_{\substack{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \leq \sigma\sqrt{\frac{m}{2}} \\ \mathbf{Az} = \mathbf{y} \pmod q}} \rho_{\frac{\sigma}{\sqrt{2}}}(\mathbf{z})} \right\} \leq 2^{-\Omega(m)}.$$

Second, we can replace the procedure $\mathsf{Revoke}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \rho_R)$ by an (inefficient) projective measurement $\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, I - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\}$, since they produce statistically close outcomes. This follows from the fact that $\mathsf{Revoke}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \rho_{\text{R}})$ applies the procedure $\mathsf{QSampGauss}$ in Algorithm 2 as a subroutine, which is correct with overwhelming probability acccording to Theorem 3.3.

Let us now analyze the success probability of Algorithm 3. Putting everything together, we get

$$
\begin{aligned}
&\Pr\left[\begin{array}{c}\mathbf{x}\leftarrow\mathsf{SIS\_Solver}(\mathbf{A})\\ \bigwedge\\ \mathbf{x}\neq\mathbf{0}\ \text{s.t.}\ \|\mathbf{x}\|\leq\sigma\sqrt{2m}\end{array} \;:\; \mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\right]\\
&\geq\left(1-\max_{\substack{\mathbf{x}\in\mathbb{Z}_q^m,\,\|\mathbf{x}\|\leq\sigma\sqrt{\frac{m}{2}}\\ \mathbf{A}\mathbf{x}=\mathbf{y}\ (\mathrm{mod}\ q)}}\left\{\frac{\rho_{\frac{\sigma}{\sqrt{2}}}(\mathbf{x})}{\sum_{\substack{\mathbf{z}\in\mathbb{Z}_q^m,\|\mathbf{z}\|\leq\sigma\sqrt{\frac{m}{2}}\\ \mathbf{A}\mathbf{z}=\mathbf{y}\ (\mathrm{mod}\ q)}}\rho_{\frac{\sigma}{\sqrt{2}}}(\mathbf{z})}\right\}\right)\\
&\quad\cdot\Pr\left[\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\rho_R)=\top\;:\;\begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\ \text{s.t.}\ \mathbf{A}\ \text{is full-rank}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\mathrm{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|)\end{array}\right]\\
&\quad\cdot\Pr\left[\mathcal{E}\left(\mathbf{A},\mathbf{y},\rho_{\mathrm{Aux}}\right)\;\in\;\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{m/2})\;:\;\begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\ \text{s.t.}\ \mathbf{A}\ \text{is full-rank}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\mathrm{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|)\\ \top\leftarrow\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\rho_{\mathrm{R}})\end{array}\right]\\
&\geq\left(1-2^{-\Omega(m)}\right)\cdot\Pr\left[\begin{array}{c}\mathsf{IneffRevoke}(\mathbf{A},\mathbf{y},\sigma,\mathrm{R})=\top\\ \bigwedge\\ \mathcal{E}(\mathbf{A},\mathbf{y},\mathrm{Aux})\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array}\;:\;\begin{array}{c}\mathbf{A}\xleftarrow{\$}\mathbb{Z}_q^{n\times m}\ \text{s.t.}\ \mathbf{A}\ \text{is full-rank}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \rho_{\mathrm{R,Aux}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|)\end{array}\right]\\
&\geq\left(1-2^{-\Omega(m)}\right)\cdot\left(\mathrm{poly}(\varepsilon,1/q)-q^{-n}\right)\quad\geq\quad\mathrm{poly}(\varepsilon,1/q).
\end{aligned}
$$

In the last line, we applied the simultaneous search-to-decision reduction from Theorem 6.19 and Lemma 2.4. Therefore, $\mathsf{SIS\_Solver}(\mathbf{A})$ in Algorithm 3 runs in time $\mathrm{poly}(q,1/\varepsilon)$ and solves $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ whenever $\varepsilon=1/\mathrm{poly}(\lambda)$. Therefore, we conclude that $\varepsilon(\lambda)$ must be negligible. □

6.6 Proof of Theorem 6.2

Proof. The proof is the same as in Theorem 6.1, except that we invoke Theorem 6.16 instead of Theorem 6.19 to argue that simultaneous extraction succeeds with sufficiently high probability. □

7 Key-Revocable Dual-Regev Encryption with Classical Revocation

Recall that our key-revocable Dual-Regev public-key encryption scheme in Construction 1 requires that a quantum state is *returned* as part of revocation. In this section, we present a Dual-Regev public-key encryption scheme with *classical key revocation* (Construction 2).

To prove security, we follow a similar proof as in our previous construction from Section 6. As an additional technical ingredient, we rely on a strong type of *collapsing* property Ajtai hash recently proven by Bartusek, Khurana and Poremba [BKP23, Theorem 5.5], which relies on the hardness of subexponential LWE and SIS.

7.1 Definition: Public-Key Encryption with Classical Key Revocation

Let us now give a formal definition of a public-key encryption scheme with classical key revocation.

Definition 7.1 (Key-Revocable Public-Key Encryption with Classical Revocation)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. A key-revocable public-key encryption scheme with classical revocation consists efficient algorithms* $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Delete}, \mathsf{Revoke})$*, where* Enc *and* Revoke *are* PPT *algorithms, and* $\mathsf{KeyGen}, \mathsf{Dec}$ *and* Delete *are* QPT *algorithms defined as follows:*

- $\mathsf{KeyGen}(1^\lambda)$*: given as input a security parameter λ, output a public key* PK*, a master secret key* MSK *and a quantum decryption key* ρ_{SK}*.*
- $\mathsf{Enc}(\mathsf{PK}, x)$*: given a public key* PK *and plaintext $x \in \{0,1\}^\ell$, output a ciphertext* CT*.*
- $\mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT})$*: given a decryption key* ρ_{SK} *and ciphertext* CT*, output a message y.*
- $\mathsf{Delete}(\rho_{\mathsf{SK}})$*: given a quantum decryption key, it outputs a classical certificate π.*
- $\mathsf{Revoke}(\mathsf{PK}, \mathsf{MSK}, \pi)$*: given as input a master secret key* MSK*, a public key* PK *and a certificate π, output* Valid *or* $\mathsf{Invalid}$*.*

Correctness of Decryption. For every $x \in \{0,1\}^\ell$, the following holds:

$$\Pr\left[x \leftarrow \mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT}) \ : \ \begin{smallmatrix}(\mathsf{PK},\mathsf{MSK},\rho_{\mathsf{SK}}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{PK}, x)\end{smallmatrix}\right] \geq 1 - \nu(\lambda),$$

where $\nu(\cdot)$ is a negligible function.

Correctness of Revocation. The following holds:

$$\Pr\left[\mathsf{Valid} \leftarrow \mathsf{Revoke}(\mathsf{PK}, \mathsf{MSK}, \pi) \ : \ \begin{smallmatrix}(\mathsf{PK},\mathsf{MSK},\rho_{\mathsf{SK}}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \pi \leftarrow \mathsf{Delete}(\rho_{\mathsf{SK}})\end{smallmatrix}\right] \geq 1 - \nu(\lambda),$$

where $\nu(\cdot)$ is a negligible function.

Our security definition for key-revocable public-key encryption is as follows.

$\underline{\mathsf{Expt}_{\Sigma,\mathcal{A}}\left(1^\lambda, b\right)}$:

Initialization Phase:

- The challenger runs $(\mathsf{PK}, \mathsf{MSK}, \rho_{\mathsf{SK}}) \leftarrow \mathsf{KeyGen}(1^\lambda)$ and sends $(\mathsf{PK}, \rho_{\mathsf{SK}})$ to $\mathcal{A}$.

Revocation Phase:

- The challenger sends the message `REVOKE` to $\mathcal{A}$.
- The adversary $\mathcal{A}$ returns a certificate π.
- The challenger aborts if $\mathsf{Revoke}(\mathsf{PK}, \mathsf{MSK}, \pi)$ outputs $\mathsf{Invalid}$.

Guessing Phase:

- $\mathcal{A}$ submits a plaintext $x \in \{0,1\}^\ell$ to the challenger.
- If $b = 0$: The challenger sends $\mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{PK}, x)$ to $\mathcal{A}$. Else, if $b = 1$, the challenger sends $\mathsf{CT} \xleftarrow{\$} \mathcal{C}$, where $\mathcal{C}$ is the ciphertext space of ℓ bit messages.
- Output $b_{\mathcal{A}}$ if the output of $\mathcal{A}$ is $b_{\mathcal{A}}$.

Figure 14: Security Experiment for Classical Key Revocation

Definition 7.2. *A public-key encryption scheme* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Delete}, \mathsf{Revoke})$ *with classical key revocation is* (ε, δ)*-secure if, for every* QPT *adversary* $\mathcal{A}$ *with*

$$\Pr[\mathsf{Invalid} \leftarrow \mathsf{Expt}_{\Sigma,\mathcal{A}}(1^\lambda, b)] \leq \delta(\lambda)$$

for $b \in \{0,1\}$*, it holds that*

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}_{\Sigma,\mathcal{A}}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}_{\Sigma,\mathcal{A}}(1^\lambda, 1)\right]\right| \leq \varepsilon(\lambda),$$

where $\mathsf{Expt}_{\Sigma,\mathcal{A}}(1^\lambda, b)$ *is as defined in Figure 14. If* $\delta(\lambda) = 1 - 1/\mathrm{poly}(\lambda)$ *and* $\varepsilon(\lambda) = \mathsf{negl}(\lambda)$*, we simply say the key-revocable public-key encryption scheme is secure.*

7.2 Construction.

Preliminaries. Let us first introduce some additional background on lattice trapdoors.

Definition 7.3 (β-good trapdoor)**.** *Let* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ *be a matrix. A* β*-good lattice trapdoor* $\mathbf{T} \in \mathbb{Z}^{m \times m}$ *for the matrix* $\mathbf{A}$ *is a short basis for the lattice* $\Lambda_q^\perp(\mathbf{A})$ *with the following properties:*

1. *Each column vector of* $\mathbf{T}$ *is in the right kernel of* $\mathbf{A}$ *(mod* q*), i.e.* $\mathbf{A} \cdot \mathbf{T} = \mathbf{0} \pmod q$.

2. *Each column vector of* $\mathbf{T} = [\mathbf{t}_1, \ldots, \mathbf{t}_m]$ *is short, i.e.* $\|\mathbf{t}_i\| \leq \beta$, *for all* $i \in [m]$.

3. *The matrix* $\mathbf{T}$ *has full rank* m *over* $\mathbb{R}$.

We make use of the following specific variant of a lattice trapdoor for the q-ary lattice $\Lambda_q^{\perp}(\mathbf{A})$ due to Alwen and Peikert [AP08].

Lemma 7.4 (AP lattice trapdoor, [AP08], Lemma 3.5)**.** *Let* $n, m \in \mathbb{N}$ *and* $q \in \mathbb{N}$ *be a prime such that* $m = \lceil 6n \log q \rceil$. *There exists a randomized algorithm with the following properties:*

- $\mathsf{GenTrap}(1^n, 1^m, q)$*: on input* $1^n, 1^m$ *and* q, *returns a matrix* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ *and a short trapdoor matrix* $\mathbf{T}_{\mathbf{A}} \in \mathbb{Z}^{m \times m}$ *such that the distribution of* $\mathbf{A}$ *is negligibly (in the parameter* n*) close to uniform and* $\mathbf{T}_{\mathbf{A}}$ *is a* $(20n \log q)$*-good lattice trapdoor with all but negligible probability.*

Let us now present our Dual-Regev encryption scheme with classical key revocation

Parameters. Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n, m \in \mathbb{N}$ and $q \in \mathbb{N}$ be a prime with $q = 2^{o(n)}$ such that $m = \lceil 6n \log q \rceil$ and $\nu = 64m^2$, each parameterized by λ. Choose a parameter $\sigma \in \mathbb{R}$ such that $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and $\sigma/\nu = 2^{o(n)}$. Moreover, let $\alpha, \beta \in (0,1)$ be noise ratios with the property that $\beta/\alpha = 2^{o(n)}$ such that $1/\alpha = 2^{o(n)} \cdot \sigma$.

Construction 2 (Dual-Regev Encryption with Classical Key Revocation)**.** *Let* $\lambda \in \mathbb{N}$ *be the security parameter. The key-revocable scheme* $\mathsf{CRevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Delete}, \mathsf{Revoke})$ *with classical revocation consists of the following* QPT *algorithms:*

- $\mathsf{KeyGen}(1^\lambda) \to (\mathsf{PK}, \rho_{\mathsf{SK}}, \mathsf{MSK})$: *sample* $(\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{T}_{\mathbf{A}} \in \mathbb{Z}^{m \times m}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ *using the procedure in Lemma 7.4, sample a random vector* $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$, *generate a Gaussian superposition* $(|\psi_{\mathbf{y}}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$, *and let*

$$|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle = \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} = \mathbf{y}}} \rho_\sigma(\mathbf{x})\, \omega_q^{\langle \mathbf{x}, \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v} \rangle} |\mathbf{x}\rangle .$$

  Output $\mathsf{PK} = (\mathbf{A}, \mathbf{y})$, $\rho_{\mathsf{SK}} = |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle$ *and* $\mathsf{MSK} = (\mathbf{T}_{\mathbf{A}}, \mathbf{v})$.

- $\mathsf{Enc}(\mathsf{PK}, \mu) \to \mathsf{CT}$: *to encrypt a bit* $\mu \in \{0,1\}$, *sample a random vector* $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ *and errors* $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$ *and* $e' \sim D_{\mathbb{Z}, \beta q}$, *and output the ciphertext pair*

$$\mathsf{CT} = \left(\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q, \mathbf{s}^\intercal \mathbf{y} + e' + \mu \cdot \lfloor \frac{q}{2} \rfloor \pmod q \right) \in \mathbb{Z}_q^m \times \mathbb{Z}_q.$$

- $\mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT}) \to \{0,1\}$: *to decrypt* CT, *apply the unitary* $U : |\mathbf{x}\rangle |0\rangle \to |\mathbf{x}\rangle |\mathsf{CT} \cdot (-\mathbf{x}, 1)^\intercal\rangle$ *on input* $|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle |0\rangle$, *where* $\rho_{\mathsf{SK}} = |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle$, *and measure the second register in the computational basis. Output* 0, *if the measurement outcome is closer to* 0 *than to* $\lfloor \frac{q}{2} \rfloor$, *and output* 1, *otherwise.*[14]

- $\mathsf{Delete}(\rho) \to \mathbf{w}$*: apply the* m*-qudit* q*-ary quantum Fourier transform to the state* ρ, *and measure in the computational basis to obtain an outcome* $\mathbf{w} \in \mathbb{Z}_q^m$.

- $\mathsf{Revoke}(\mathsf{MSK}, \mathbf{w}) \to \{\top, \bot\}$*: to verify the certificate* $\mathbf{w} \in \mathbb{Z}_q^m$, *do the following:*

[14] This procedure can be purified via the Gentle Measurement Lemma.

1. *Parse* $(\mathbf{T_A} \in \mathbb{Z}^{m\times m}, \mathbf{v} \in \{0,1\}^m) \leftarrow \mathsf{MSK}$.
2. *Compute* $\mathbf{c}^\intercal = \mathbf{w}^\intercal \cdot \mathbf{T_A} \pmod q$.
3. *Compute* $\mathbf{d}^\intercal = \mathbf{c}^\intercal \cdot \mathbf{T_A^{-1}}$, *where* $\mathbf{T_A^{-1}}$ *is the inverse matrix of* $\mathbf{T_A}$ *over* $\mathbb{R}$.[15]
4. *For* $i \in [m]$, *let* $v'_i = 0$, *if* $d_i \in [-q/\sigma, q/\sigma]$, *and let* $v'_i = 1$, *otherwise.*
5. *Output* $\top$, *if* $\mathbf{v}' = \mathbf{v}$, *and output* $\bot$ *otherwise.*

Correctness of verification.

Lemma 7.5. *Let* $\lambda \in \mathbb{N}$ *be the security parameter. Then, by our choice of parameters.*

$$\mathsf{Pr}\left[\top \leftarrow \mathsf{Revoke}(\mathsf{MSK}, \mathbf{w}) \ : \ \begin{smallmatrix}(\mathsf{PK},\mathsf{MSK},\rho_{\mathsf{SK}})\leftarrow\mathsf{KeyGen}(1^\lambda)\\ \mathbf{w}\leftarrow\mathsf{Delete}(\rho_{\mathsf{SK}})\end{smallmatrix}\right] \geq 1 - \mathsf{negl}(\lambda)$$

Proof. Recall that $\mathsf{KeyGen}(1^\lambda)$ samples $(\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathbf{T} \in \mathbb{Z}^{m\times m}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ and $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$, generates a Gaussian superposition $(|\psi_{\mathbf{y}}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ and assigns $\mathsf{PK} = (\mathbf{A}, \mathbf{y})$, $\rho_{\mathsf{SK}} = |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle$ and $\mathsf{MSK} = (\mathbf{T}, \mathbf{v})$, where

$$|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle = \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{Ax}=\mathbf{y}}} \rho_\sigma(\mathbf{x})\, \omega_q^{\langle \mathbf{x}, \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}\rangle} |\mathbf{x}\rangle .$$

From Lemma 7.4 and Lemma 2.4, it follows that $\mathbf{A}$ is full rank and that $\mathbf{T}$ is a $(20n \log q)$-good lattice trapdoor with overwhelming probability. Moreover, from Lemma 3.2, it follows that $\mathbf{w} \leftarrow \mathsf{Del}(\rho_{\mathsf{SK}})$ results in a vector $\mathbf{w}^\intercal = \hat{\mathbf{s}}^\intercal \mathbf{A} + \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}^\intercal + \hat{\mathbf{e}}^\intercal \in \mathbb{Z}_q^m$ with $\hat{\mathbf{e}} \sim D_{\mathbb{Z}_q^m, q/\sigma}$. Therefore, we have that with overwhelming probability $\|\hat{\mathbf{e}}\|_2 \leq \sqrt{m}q/\sigma$ and thus

$$\begin{aligned}
\|\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}^\intercal \cdot \mathbf{T} + \hat{\mathbf{e}}^\intercal \cdot \mathbf{T}\|_\infty &\leq \lfloor \frac{q}{\nu} \rfloor \cdot \|\mathbf{v}^\intercal \mathbf{T}\|_2 + \|\hat{\mathbf{e}}^\intercal \mathbf{T}\|_2 \\
&\leq \lfloor \frac{q}{\nu} \rfloor \cdot \|\mathbf{v}\|_2 \cdot \|\mathbf{T}\|_2 + \|\hat{\mathbf{e}}\|_2 \cdot \|\mathbf{T}\|_2 \\
&\leq \lfloor \frac{q}{\nu} \rfloor \cdot \sqrt{m} \cdot \|\mathbf{T}\|_2 + \sqrt{m} \cdot \frac{q}{\sigma} \cdot \|\mathbf{T}\|_2 \\
&\leq \left(\lfloor \frac{q}{\nu} \rfloor + \frac{q}{\sigma}\right) \cdot m \cdot \|\mathbf{T}\| \ \leq \ q/4.
\end{aligned}$$

Consequently, $\mathbf{c}^\intercal = \mathbf{w}^\intercal \cdot \mathbf{T} = \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}^\intercal \cdot \mathbf{T} + \hat{\mathbf{e}}^\intercal \cdot \mathbf{T} \in \mathbb{Z}^m \cap (-\frac{q}{4}, \frac{q}{4}]^m$ and thus

$$\mathbf{d}^\intercal = \mathbf{c}^\intercal \cdot \mathbf{T}^{-1} = \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}^\intercal + \hat{\mathbf{e}}^\intercal.$$

By our choice of parameters, we have that $\lfloor \frac{q}{\nu} \rfloor \gg \frac{q}{\sigma}$, and thus the procedure $\mathsf{Revoke}(\mathsf{MSK}, \mathbf{w})$ successfully decodes $\mathbf{v}$ from $\mathbf{d}$, and outputs $\top$ with overwhelming probability.

□

7.3 Proof of Security.

We first discuss some relevant background which we use for the security proof.

[15] The matrix $\mathbf{T_A} \in \mathbb{Z}^{m\times m}$ output by $\mathsf{GenTrap}(1^n, 1^m, q)$ has full rank m over $\mathbb{R}$ with overwhelming probability.

Everlasting collapsing property of the Ajtai hash function. We need the following result from Bartusek, Khurana and Poremba [BKP23, Theorem 5.5]. Informally, it says that once the adversary outputs a short pre-image (given as input either a Gaussian superposition of pre-images, or a single measured pre-image), then the adversary cannot tell which state it initially received even if it is allowed to be unbounded for the rest of the experiment. This property is called *certified-everlasting target-collapsing*[16] and it can be shown assuming the hardness of LWE/SIS.

Theorem 7.6 (Certified-everlasting target-collapsing property of the Ajtai hash, [BKP23])**.** *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $q/\sqrt{8m} > \sigma > \sqrt{8m}$ and let $\alpha \in (0,1)$ be a noise ratio such that $1/\alpha = 2^{o(n)} \cdot \sigma$. Then, assuming the hardness of the $\mathsf{LWE}^m_{n,q,\alpha q}$ and $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$, it holds for any adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ consisting of a QPT algorithm $\mathcal{A}_0$ and an unbounded algorithm $\mathcal{A}_1$:*

$$|\Pr\left[\mathsf{EvTargetCollapseExpt}_{\mathcal{A},\lambda}(0) = 1\right] - \Pr\left[\mathsf{EvTargetCollapseExpt}_{\mathcal{A},\lambda}(1) = 1\right]| \leq \mathsf{negl}(\lambda).$$

Here, the experiment $\mathsf{EvTargetCollapseExpt}_{\mathcal{A},\lambda}(b)$ is defined as follows:

1. *The challenger generates a Gaussian superposition $(|\psi_{\mathbf{y}}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$, with*

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} = \mathbf{y}}} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \, .$$

2. *If $b = 0$, the challenger does nothing. Else, if $b = 1$, the challenger measures the state in the computational basis. Next, the challenger sends the resulting state together with y to $\mathcal{A}_0$.*

3. *$\mathcal{A}_0$ sends a classical certificate $\mathbf{w} \in \mathbb{Z}_q^m$ to the challenger and initializes $\mathcal{A}_1$ with its residual (internal) state.*

4. *The challenger checks if $\mathbf{w}$ satisfies $\mathbf{A} \cdot \mathbf{w} = \mathbf{y} \pmod q$ and $\|\mathbf{w}\| \leq \sigma\sqrt{m/2}$. If true, $\mathcal{A}_1$ is run until it outputs a bit b'. Otherwise, $b' \leftarrow \{0,1\}$ is sampled uniformly at random. The output of the experiment is b'.*

We immediately obtain result.

Theorem 7.7 (Uncertainty principle for Gaussian states)**.** *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $q/\sqrt{8m} > \sigma > \sqrt{8m}$ and let $\alpha \in (0,1)$ be a noise ratio such that $1/\alpha = 2^{o(n)} \cdot \sigma$, and let $\nu > 0$. Then, assuming the subexponential hardness of the $\mathsf{LWE}^m_{n,q,\alpha q}$, it holds for any adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ consisting of a QPT algorithm $\mathcal{A}_0$ and an unbounded algorithm $\mathcal{A}_1$:*

$$\Pr[\mathsf{UncertaintyExpt}_{\mathcal{A}}(1^\lambda) = 1] \leq \mathsf{negl}(\lambda),$$

where $\mathsf{UncertaintyExpt}_{\mathcal{A}}(1^\lambda)$ is the following experiment between a challenger and the adversary $\mathcal{A}$:

[16] It's called "everlasting" because the adversary can be unbounded once a pre-image is presented and it's called "target-collapsing" because it is a weakening of collapsing, where the image is honestly generated by the challenger.

1. *The challenger samples a random matrix* $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ *and vector* $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$*, generates a Gaussian superposition* $(|\psi_{\mathbf{y}}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ *and prepares*

$$|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle = \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y}}} \rho_\sigma(\mathbf{x})\, \omega_q^{\langle \mathbf{x}, \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}\rangle} |\mathbf{x}\rangle .$$

   Then, the challenger sends $(\mathbf{A}, \mathbf{y}, |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle)$ *to* $\mathcal{A}_0$*.*

2. $\mathcal{A}_0$ *generates a vector* $\mathbf{x}_0 \in \mathbb{Z}_q^m$ *and a polynomial-sized quantum state* ρ_{AUX}*. Then,* $\mathcal{A}_0$ *sends* $\mathbf{x}_0$ *to the challenger and forwards* $(\mathbf{A}, \mathbf{y}, \rho_{\text{AUX}})$ *to* $\mathcal{A}_1$*.*

3. $\mathcal{A}_1$ *generates a vector* $\mathbf{v}' \in \mathbb{Z}_q^m$ *on input* $(\mathbf{A}, \mathbf{y}, \rho_{\text{AUX}})$ *and sends it to the challenger.*

4. *The challenger checks if* $\mathbf{v}' = \mathbf{v}$ *and whether the vector* $\mathbf{x}_0$ *satisfies* $\mathbf{A} \cdot \mathbf{x}_0 = \mathbf{y} \pmod q$ *and* $\|\mathbf{x}_0\| \leq \sigma\sqrt{m/2}$*. If true, the challenger outputs* 1*, otherwise the challenger outputs* 0*. The bit output by the challenger also denotes the outcome of the experiment.*

Proof. Let $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ be an adversary consisting of a QPT algorithm $\mathcal{A}_0$ and an unbounded algorithm $\mathcal{A}_1$. We consider the following two hybrids:

H_0: This is the original experiment $\mathsf{UncertaintyExpt}_{\mathcal{A}}(1^\lambda)$.

H_1: This is the same experiment as before, except that the challenger additionally measures the state $|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle$ in Step 1 before sending it to $\mathcal{A}_0$.

To complete the proof, we show the following:

Claim 7.8.

$$|\Pr[\mathsf{H}_0 = 1] - \Pr[\mathsf{H}_1 = 1]| \leq \mathsf{negl}(\lambda).$$

Proof. This follows from the certified-everlasting target-collapsing property of the Ajtai hash in Theorem 7.6. Suppose for the sake of contradiction that the advantage is $1/\mathrm{poly}(\lambda)$. We consider the following adversary $\mathcal{B} = (\mathcal{B}_0, \mathcal{B}_1)$ consisting of a QPT algorithm $\mathcal{B}_0$ and an unbounded algorithm $\mathcal{B}_1$ that distinguishes $\mathsf{EvTargetCollapseExpt}_{\mathcal{B},\lambda}(1^\lambda, b)$ for $b \in \{0,1\}$ with advantage $1/\mathrm{poly}(\lambda)$.

- $\mathcal{B}_0$ receives as input $(\mathbf{A}, \mathbf{y}, |\psi_{\mathbf{y}}\rangle)$ from the challenger, samples $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$, generates the state $|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle = \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{y}}\rangle$ and then forwards $(\mathbf{A}, \mathbf{y}, |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle)$ to $\mathcal{A}_0$.
- When $\mathcal{A}_0$ outputs a vector $\mathbf{x}_0 \in \mathbb{Z}_q^m$ and a polynomial-size advice state ρ_{AUX}, $\mathcal{B}_0$ forwards $\mathbf{x}_0$ to the challenger, and initializes $\mathcal{B}_1$ with ρ_{AUX}.
- $\mathcal{B}_1$ runs $\mathcal{A}_1$ with auxuliary input ρ_{AUX}. When $\mathcal{A}_1$ outputs a vector $\mathbf{v}'$, the procedure $\mathcal{B}_1$ checks if $\mathbf{v}' = \mathbf{v}$. If true, $\mathcal{B}_1$ outputs 1. Else, $\mathcal{B}_1$ outputs 0 otherwise.

Note that if $b = 0$, the state $|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle$ prepared by $\mathcal{B}_0$ corresponds to

$$|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y}}} \rho_\sigma(\mathbf{x})\, \omega_q^{\langle \mathbf{x}, \lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}\rangle} |\mathbf{x}\rangle ,$$

and thus matches the correct input from H_0. Moreover, if $b = 1$, the state is collapsed and thus matches the correct input from H_1. Crucially, the Pauli-Z operator $\mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}}$ has no effect on the collapsed state and merely introduces a global phase. Therefore, $\mathcal{B} = (\mathcal{B}_0, \mathcal{B}_1)$ distinguishes $\mathsf{EvTargetCollapseExpt}_{\mathcal{B},\lambda}(1^\lambda, b)$ for $b \in \{0,1\}$ with advantage at least $1/\mathrm{poly}(\lambda)$, which results in the desired contradiction. □

Because the vector $\mathbf{v}$ is completely hidden from the view of the adversary in H_1, the probability that $\mathcal{A}_1$ guesses it correctly is at most 2^{-m}, which is negligible. This proves the claim. □

Simultaneous search-to-decision reduction with classical revocation We give a strengthening of our result in Theorem 6.11 and state a *simultaneous* search-to-decision reduction with quantum auxiliary input which holds even if additionally require that a *classical revocation* procedure succeeds. This is essentially a classical revocation analogue of Theorem 6.16.

To prove the theorem, we require the following analogue of Conjecture 1.

Conjecture 2. *Let $\lambda \in \mathbb{N}$. Then, there exist parameters (each parameterized by λ) such that $n \in \mathbb{N}$, q is a prime with $q = 2^{o(n)}$, $m \geq 2n \log q$, $\sqrt{8m} < \sigma < q/\sqrt{8m}$, $\alpha \in (0,1)$ with $1/\alpha = 2^{o(n)} \cdot \sigma$ and $\nu = 64m^2$ for which the following holds: for any QPT $\mathcal{A}$ and QPT $\mathcal{C}$ (and for a fixed algorithm $\mathcal{B}$), and for any $\mathrm{poly}(\lambda)$-sized quantum auxiliary input τ_λ (which depends on λ):*

$$\left|\Pr[1 \leftarrow \mathsf{Expt}_{\mathcal{A},\mathcal{B},\mathcal{C}}(1^\lambda, 0)] - \Pr[1 \leftarrow \mathsf{Expt}_{\mathcal{A},\mathcal{B},\mathcal{C}}(1^\lambda, 1)]\right| \leq \mathsf{negl}(\lambda),$$

where $\mathsf{Expt}_{\mathcal{A},\mathcal{B},\mathcal{C}}(1^\lambda, b)$ is the cloning experiment in Figure 8.

We prove the following statement.

Theorem 7.9 (Simultaneous Search-to-Decision Reduction with Classical Revocation)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n \in \mathbb{N}$, q be a prime with $q = 2^{o(n)}$, $m \geq 6n \log q$, $\sqrt{8m} < \sigma < q/\sqrt{8m}$, $\alpha, \beta \in (0,1)$ with $\beta/\alpha = 2^{o(n)}$, $1/\alpha = 2^{o(n)} \cdot \sigma$ and $\nu = 64m^2$, such that the following holds: Let $\mathcal{A} = \{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}, \tau_\lambda)\}_{\lambda \in \mathbb{N}}$ be any non-uniform quantum algorithm consisting of a family of polynomial-sized quantum circuits*

$$\left\{\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}} : \mathcal{L}(\mathcal{H}_q^m \otimes \mathcal{H}_{B_\lambda}) \to \mathcal{L}(\mathcal{H}_{R_\lambda} \otimes \mathcal{H}_{\mathrm{AUX}_\lambda})\right\}_{\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{y} \in \mathbb{Z}_q^n}$$

and polynomial-sized advice states $\tau_\lambda \in \mathcal{D}(\mathcal{H}_{B_\lambda})$ which are independent of $\mathbf{A}$. Then, assuming Conjecture 2 is true, the following statement holds for every QPT distinguisher $\mathcal{D}$. Suppose there exists a function $\varepsilon(\lambda) = 1/\mathrm{poly}(\lambda)$ such that

$$\Big|\Pr\left[1 \leftarrow \mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda, 1)\right]\Big| = \varepsilon(\lambda).$$

Then, there exists a quantum extractor $\mathcal{E}$ that takes as input $\mathbf{A}$, $\mathbf{y}$ and system AUX of the state $\rho_{\mathrm{R,AUX}}$ and outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\mathrm{poly}(\lambda, m, \sigma, q, 1/\varepsilon)$ such that

$$\Pr\left[\begin{array}{c} \mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top \\ \wedge \\ \mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}}) \in \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{\frac{m}{2}}) \end{array} : \begin{array}{c} (\mathbf{A},\mathbf{T_A}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q) \\ (|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma) \\ \mathbf{v} \xleftarrow{\$} \{0,1\}^m,\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle \leftarrow \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{y}}\rangle \\ |\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\mathrm{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}| \otimes \tau_\lambda) \end{array}\right] \geq 1/\mathrm{poly}(\varepsilon, 1/q).$$

$\underline{\mathsf{Expt}_{\mathcal{A},\mathcal{B},\mathcal{C}}(1^\lambda, b)}$:

1. The challenger samples $(\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathbf{T_A}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ and a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and lets $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}_0 \pmod q$. Then, the challenger runs $|\psi_{\mathbf{y}}\rangle \leftarrow \mathsf{QSampGauss}(\mathbf{A}, \mathbf{T_A}, \mathbf{y}, \sigma)$, samples $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$ and sends $(\mathbf{A}, \mathbf{y}, |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle)$ to $\mathcal{A}$, where $|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle = \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{y}}\rangle$.
2. $\mathcal{A}$ receives $(\mathbf{A}, \mathbf{y}, |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle)$ together with auxiliary input τ_λ and generates a state ρ_{BC} in systems BC where B is classical, and sends B to $\mathcal{B}$ and C to $\mathcal{C}$.
3. The challenger sends $(\mathbf{T_A}, \mathbf{v}, \sigma)$ to $\mathcal{B}$ and, depending on the value of b, the challenger sends the following to $\mathcal{C}$:
   - if $b = 0$: the challenger samples $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$, lets $\mathbf{u} = \mathbf{A}^\intercal \mathbf{s} + \mathbf{e}$, and sends $(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal \mathbf{x}_0 \pmod q)$ to $\mathcal{C}$.
   - if $b = 0$: the challenger samples a uniformly random vector $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and sends $(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal \mathbf{x}_0 \pmod q)$ to $\mathcal{C}$.
4. $\mathcal{B}$ receives as input $(\mathbf{T_A}, \mathbf{v}, \sigma, \mathbf{w})$, where $\mathbf{w}$ is the content of the classical register B, and applies $\mathsf{Revoke}(\mathbf{T_A}, \mathbf{v}, \mathbf{w})$. $\mathcal{B}$ outputs $\top$, if it succeeds, else outputs $\bot$.
5. $\mathcal{C}$ receives as input $(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal \mathbf{x}_0 \pmod q, \mathrm{C})$ and outputs a bit b'.
6. The challenger outputs 1, if $\mathcal{B}$ outputs $\top$ and $\mathcal{C}$ outputs $b' = b$. This is also the outcome of the experiment.

Figure 15: The experiment for Conjecture 2.

$\underline{\mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}\left(1^\lambda, b\right)}$:

- If $b = 0$: output $\mathsf{simult.lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 17.
- If $b = 1$: output $\mathsf{simult.unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 18.

Figure 16: The experiment $\mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}\left(1^\lambda, b\right)$.

$\underline{\mathsf{simult.lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)}$:

1. Sample $(\mathbf{A}, \mathbf{T}_\mathbf{A}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$.
2. Generate $(|\psi_\mathbf{y}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$.
3. Sample $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$ and compute $|\psi^\mathbf{v}_\mathbf{y}\rangle \leftarrow \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_\mathbf{y}\rangle$.
4. Generate $|\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\mathrm{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi^\mathbf{v}_\mathbf{y}\rangle\langle\psi^\mathbf{v}_\mathbf{y}| \otimes \tau_\lambda)$.
5. Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n, \mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ and $e' \sim D_{\mathbb{Z},\beta q}$.
6. Run $\mathsf{Revoke}(\mathbf{T}_\mathbf{A}, \mathbf{w})$ on $\mathbf{w}$. If it outputs $\top$, continue. Otherwise, output $\mathsf{Invalid}$.
7. Run $b' \leftarrow \mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal, \mathbf{s}^\intercal\mathbf{y} + e', \rho_{\mathrm{AUX}})$. Output b'.

Figure 17: The distribution $\mathsf{simult.lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$.

Proof. Let $\lambda \in \mathbb{N}$ be the security parameter and let $\mathcal{A} = \{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}, \tau_\lambda)\}_{\mathbf{A} \in \mathbb{Z}_q^{n \times m}}$ be a non-uniform quantum algorithm. Suppose that $\mathcal{D}$ is a QPT distinguisher with advantage $\varepsilon = 1/\mathrm{poly}(\lambda)$.

To prove the claim, we consider the following sequence of hybrid distributions.

H_0: This is the distribution $\mathsf{simult.lwe.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ in Figure 17.

H_1: This is the following distribution:

1. Sample $(\mathbf{A}, \mathbf{T}_\mathbf{A}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$.
2. Sample a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and let $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}_0 \pmod q$.
3. Run $|\psi_\mathbf{y}\rangle \leftarrow \mathsf{QSampGauss}(\mathbf{A}, \mathbf{T}_\mathbf{A}, \mathbf{y}, \sigma)$.
4. Sample $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$ and compute $|\psi^\mathbf{v}_\mathbf{y}\rangle \leftarrow \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_\mathbf{y}\rangle$.

$\underline{\mathsf{simult.unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)}$:

1. Sample $(\mathbf{A}, \mathbf{T_A}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$.
2. Generate $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$.
3. Sample $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$ and compute $|\psi^{\mathbf{v}}_{\mathbf{y}}\rangle \leftarrow \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{y}}\rangle$.
4. Generate $|\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\text{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi^{\mathbf{v}}_{\mathbf{y}}\rangle\langle\psi^{\mathbf{v}}_{\mathbf{y}}| \otimes \tau_\lambda)$.
5. Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n, \mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ and $e' \sim D_{\mathbb{Z},\beta q}$.
6. Sample $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and $r \xleftarrow{\$} \mathbb{Z}_q$.
7. Run $\mathsf{Revoke}(\mathbf{T_A}, \mathbf{w})$ on $\mathbf{w}$. If it outputs $\top$, continue. Otherwise, output $\mathsf{Invalid}$.
8. Run $b' \leftarrow \mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{u}, r, \rho_{\text{AUX}})$. Output b'.

Figure 18: The distribution $\mathsf{simult.unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$.

5. Run $\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi^{\mathbf{v}}_{\mathbf{y}}\rangle\langle\psi^{\mathbf{v}}_{\mathbf{y}}| \otimes \tau_\lambda)$ to generate a state $|\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\text{AUX}}$.
6. Run $\mathsf{Revoke}(\mathbf{T_A}, \mathbf{v}, \mathbf{w})$. If it outputs $\top$, continue. Otherwise, output $\mathsf{Invalid}$.
7. Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ and $e' \sim D_{\mathbb{Z},\beta q}$. Let $\mathbf{u} = \mathbf{A}^\intercal\mathbf{s} + \mathbf{e}$.
8. Run the distinguisher $\mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 + e', \cdot)$ on system AUX. Output b'.

H_2 : This is the following distribution:

1. Sample $(\mathbf{A}, \mathbf{T_A}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$.
2. Sample a Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and let $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}_0 \pmod q$.
3. Run $|\psi_{\mathbf{y}}\rangle \leftarrow \mathsf{QSampGauss}(\mathbf{A}, \mathbf{T_A}, \mathbf{y}, \sigma)$.
4. Sample $\mathbf{v} \xleftarrow{\$} \{0,1\}^m$ and compute $|\psi^{\mathbf{v}}_{\mathbf{y}}\rangle \leftarrow \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{y}}\rangle$.
5. Run $\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi^{\mathbf{v}}_{\mathbf{y}}\rangle\langle\psi^{\mathbf{v}}_{\mathbf{y}}| \otimes \tau_\lambda)$ to generate a state $|\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\text{AUX}}$.
6. Run $\mathsf{Revoke}(\mathbf{T_A}, \mathbf{v}, \mathbf{w})$. If it outputs $\top$, continue. Otherwise, output $\mathsf{Invalid}$.
7. Sample $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and $e' \sim D_{\mathbb{Z},\beta q}$.
8. Run the distinguisher $\mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{u}^\intercal\mathbf{x}_0 + e', \cdot)$ on system AUX. Output b'.

H_3: This is the distribution $\mathsf{simult.unif.Dist}^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 18.

We now show the following:

Claim 7.10. *Hybrids* H_0 *and* H_1 *are statistically indistinguishable. In other words,*

$$\mathsf{H}_0 \approx_s \mathsf{H}_1.$$

Proof. Here, we invoke the *noise flooding* property in Lemma 2.9 to argue that $\mathbf{e}^\intercal \mathbf{x}_0 \ll e'$ holds with overwhelming probability for our choice of parameters. Therefore, the distributions in H_0 and H_1 are computationally indistinguishable. □

Claim 7.11. *Assuming Conjecture 2 holds for our choice of parameters, the hybrids* H_1 *and* H_2 *are computationally indistinguishable,*

$$\mathsf{H}_1 \approx_c \mathsf{H}_2.$$

Proof. This follows from directly from Conjecture 2, which essentially allows us to invoke a variant of the LWE assumption, even when the procedure $\mathsf{Revoke}(\mathbf{T_A}, \mathbf{v}, \mathbf{w})$ is simultaneously taken into account. Here, we rely on the fact that during the reduction, we can simply sample $e' \sim D_{\mathbb{Z},\beta q}$ to produce an identically distributed challenge distribution. □

Recall that H_0 and H_3 can be distinguished with probability $\varepsilon = 1/\mathrm{poly}(\lambda)$. We proved that the hybrids H_0 and H_2 are computationally indistinguishable. As a consequence, it holds that hybrids H_2 and H_3 can be distinguished with probability at least $\varepsilon - \mathsf{negl}(\lambda)$.

We leverage this to obtain a Goldreich-Levin reduction. Consider the following distinguisher.

$\underline{\tilde{\mathcal{D}}\big(\mathbf{A}, \mathbf{y}, \mathbf{u}, v, \rho\big)}$:

Input: $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$, $\mathbf{u} \in \mathbb{Z}_q^n$, $v \in \mathbb{Z}_q$ and $\rho \in L(\mathcal{H}_{\mathrm{AUX}})$.
Output: A bit $b' \in \{0,1\}$.

Procedure:

1. Sample $e' \sim D_{\mathbb{Z},\beta q}$.
2. Output $b' \leftarrow \mathcal{D}\big(\mathbf{A}, \mathbf{y}, \mathbf{u}, v + e', \rho\big)$.

Figure 19: The distinguisher $\tilde{\mathcal{D}}\big(\mathbf{A}, \mathbf{y}, \mathbf{u}, v, \rho\big)$.

Note that $r + e' \pmod q$ is uniform whenever $r \xleftarrow{\$} \mathbb{Z}_q$ and $e' \sim D_{\mathbb{Z},\beta q}$. Therefore, our previous argument shows that there exists a negligible function η such that:

$$\left| \Pr\left[\begin{array}{c} \mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top \\ \wedge \\ \tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{u},\mathbf{u}^\intercal\mathbf{x}_0,\rho_{\mathrm{AUX}})=1 \end{array} : \begin{array}{c} (\mathbf{A},\mathbf{T_A}) \leftarrow \mathsf{GenTrap}(1^n,1^m,q) \\ \mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m \\ \mathbf{x}_0 \sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}}, \mathbf{y}=\mathbf{A}\mathbf{x}_0 \pmod q \\ |\psi_{\mathbf{y}}\rangle \leftarrow \mathsf{QSampGauss}(\mathbf{A},\mathbf{y},\mathbf{T_A},\sigma) \\ |\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\mathrm{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \tau_\lambda) \end{array} \right] - \Pr\left[\begin{array}{c} \mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top \\ \wedge \\ \tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{u},r,\rho_{\mathrm{AUX}})=1 \end{array} : \begin{array}{c} (\mathbf{A},\mathbf{T_A}) \leftarrow \mathsf{GenTrap}(1^n,1^m,q) \\ \mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m, r \xleftarrow{\$} \mathbb{Z}_q \\ \mathbf{x}_0 \sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}}, \mathbf{y}=\mathbf{A}\mathbf{x}_0 \pmod q \\ |\psi_{\mathbf{y}}\rangle \leftarrow \mathsf{QSampGauss}(\mathbf{A},\mathbf{y},\mathbf{T_A},\sigma) \\ |\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\mathrm{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \tau_\lambda) \end{array} \right] \right| \geq \varepsilon - \eta(\lambda).$$

From Theorem 4.3, it follows that there exists a Goldreich-Levin extractor $\mathcal{E}$ running in time $T(\mathcal{E}) = \mathrm{poly}(\lambda, n, m, \sigma, q, 1/\varepsilon)$ that outputs a short vector in $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ with probability at least

$$\Pr\left[\begin{array}{c}\mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top\\ \bigwedge\\ \mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}})\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array} : \begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ \mathbf{x}_0\sim D_{\mathbb{Z}_q^m,\frac{\sigma}{\sqrt{2}}},\mathbf{y}=\mathbf{A}\mathbf{x}_0 \pmod q\\ |\psi_{\mathbf{y}}\rangle\leftarrow\mathsf{QSampGauss}(\mathbf{A},\mathbf{y},\mathbf{T_A},\sigma)\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\otimes\tau_\lambda)\end{array}\right]\geq 1/\mathrm{poly}(\varepsilon,1/q).$$

This proves the claim. □

Towards a proof of the conjecture. We now give a proof of the *simultaneous* search-to-decision reduction (Theorem 7.9) from standard assumptions in the special case when revocation succeeds with overwhelming probability.

Theorem 7.12. *Let $n\in\mathbb{N}$. Let q be a prime with $q=2^{o(n)}$, $m\geq 6n\log q$, $\sqrt{8m}<\sigma<q/\sqrt{8m}$, and let $\alpha,\beta\in(0,1)$ with $\beta/\alpha=2^{o(n)}$ with $1/\alpha=2^{o(n)}\cdot\sigma$. Let $\nu=64m^2$. Let $\mathcal{A}=\{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}},\tau_\lambda)\}_{\lambda\in\mathbb{N}}$ be any non-uniform quantum algorithm consisting of a family of polynomial-sized quantum circuits*

$$\left\{\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}:\mathcal{L}(\mathcal{H}_q^m\otimes\mathcal{H}_{B_\lambda})\to\mathcal{L}(\mathcal{H}_{R_\lambda}\otimes\mathcal{H}_{\mathrm{AUX}_\lambda})\right\}_{\mathbf{A}\in\mathbb{Z}_q^{n\times m},\mathbf{y}\in\mathbb{Z}_q^n}$$

and polynomial-sized advice states $\tau_\lambda\in\mathcal{D}(\mathcal{H}_{B_\lambda})$ which are independent of $\mathbf{A}$ such that

$$\Pr\left[\mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top : \begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]=1-\mu(\lambda),$$

for some negligle function $\mu(\lambda)$. Then, assuming the quantum hardness of the $\mathsf{LWE}_{n,q,\alpha q}^m$ assumption, the following holds for every QPT distinguisher $\mathcal{D}$. Suppose that there exists a function $\varepsilon(\lambda)=1/\mathrm{poly}(\lambda)$ such that

$$\Big|\Pr\left[1\leftarrow\mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda,0)\right]-\Pr\left[1\leftarrow\mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}(1^\lambda,1)\right]\Big|=\varepsilon(\lambda).$$

Then, there exists a quantum extractor $\mathcal{E}$ that takes as input $\mathbf{A}$, $\mathbf{y}$ and system AUX of the state $\rho_{\mathrm{R,AUX}}$ and outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\mathrm{poly}(\lambda,m,\sigma,q,1/\varepsilon)$ such that

$$\Pr\left[\begin{array}{c}\mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})\\ \bigwedge\\ \mathcal{E}(\mathbf{A},\mathbf{y},\mathrm{AUX})\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array} : \begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]\geq\mathrm{poly}(\varepsilon,1/q).$$

Proof. Let $\lambda\in\mathbb{N}$ and suppose that there exists an adversary consisting of a pair of QPT algorithms

$(\mathcal{A},\mathcal{D})$ and a function $\varepsilon(\lambda)=1/\mathrm{poly}(\lambda)$ such that

$$\left|\Pr\left[\begin{array}{c}\mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top\\ \wedge\\ \mathcal{D}(\mathbf{A},\mathbf{y},\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal,\mathbf{s}^\intercal\mathbf{y}+e',\rho_{\mathrm{AUX}})=1\end{array}:\begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\mathbf{s}\xleftarrow{\$}\mathbb{Z}_q^n\\ \mathbf{e}\sim D_{\mathbb{Z}^m,\,\alpha q},e'\sim D_{\mathbb{Z},\,\beta q}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]-$$
$$\Pr\left[\begin{array}{c}\mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top\\ \wedge\\ \mathcal{D}(\mathbf{A},\mathbf{y},\mathbf{u},r,\rho_{\mathrm{AUX}})=1\end{array}:\begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\mathbf{u}\xleftarrow{\$}\mathbb{Z}_q^m,r\xleftarrow{\$}\mathbb{Z}_q\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]\Bigg|=\varepsilon(\lambda).$$

From Lemma 6.10, it follows that there exists a QPT distinguisher $\tilde{\mathcal{D}}$ that succeeds on the reduced state on AUX alone (where we drop the condition that $\mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})$ outputs $\top$). In other words, there exists $\bar{\epsilon}=1/\mathrm{poly}(\lambda)$ such that

$$\left|\Pr\left[\tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal,\mathbf{s}^\intercal\mathbf{y}+e',\rho_{\mathrm{AUX}})=1:\begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\mathbf{s}\xleftarrow{\$}\mathbb{Z}_q^n\\ \mathbf{e}\sim D_{\mathbb{Z}^m,\,\alpha q},e'\sim D_{\mathbb{Z},\,\beta q}\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]-$$
$$\Pr\left[\tilde{\mathcal{D}}(\mathbf{A},\mathbf{y},\mathbf{u},r,\rho_{\mathrm{AUX}})=1:\begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\mathbf{u}\xleftarrow{\$}\mathbb{Z}_q^m,r\xleftarrow{\$}\mathbb{Z}_q\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]\Bigg|=\bar{\varepsilon}(\lambda).$$

Because $\tilde{\mathcal{D}}$ succeeds with distinguishing advantage $\bar{\varepsilon}=1/\mathrm{poly}(\lambda)$, we can use Theorem 6.11 to argue that there exists a quantum extractor $\mathcal{E}$ that takes as input a set of good inputs $(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}})$, and outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\mathrm{poly}(\lambda,m,\sigma,q,1/\bar{\varepsilon})$ such that

$$\Pr\left[\begin{array}{c}\mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}})=\mathbf{x}_0\\ \wedge\\ \mathbf{x}_0\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array}:\begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\,|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]\geq\mathrm{poly}(\bar{\varepsilon},1/q).$$

Recall that revocation succeeds with overwhelming probability, i.e.,

$$\Pr\left[\mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top:\begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\,|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]=1-\mathsf{negl}(\lambda).$$

Therefore, we can use Bonferroni's inequality to argue that

$$\begin{aligned}
&\Pr\left[\begin{array}{c}\mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top\\ \bigwedge\\ \mathcal{E}(\mathbf{A},\mathbf{y},\mathrm{AUX})\in\Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}})\end{array} : \begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\mathbf{A}\rangle\langle\mathbf{A}|\otimes|\mathbf{y}\rangle\langle\mathbf{y}|\otimes|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]\\
&\geq\Pr\left[\mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top : \begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]\\
&\quad+\Pr\left[\mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}})\ \in\ \Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{m/2}) : \begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]-1\\
&\geq\Pr\left[\mathcal{E}(\mathbf{A},\mathbf{y},\rho_{\mathrm{AUX}})\ \in\ \Lambda_q^{\mathbf{y}}(\mathbf{A})\cap\mathcal{B}^m(\mathbf{0},\sigma\sqrt{m/2}) : \begin{array}{c}(\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q)\\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma)\\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle\\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda)\end{array}\right]-\mathsf{negl}(\lambda)\\
&\geq\mathrm{poly}(\varepsilon,1/q).
\end{aligned}$$

This proves the claim. □

Proof of security. We can show that the Dual-Regev scheme with classical revocation achieves our notion key-revocable security from Definition 7.2, assuming either Theorem 7.9 is true, or using Theorem 7.12 based on LWE/SIS in case revocation succeeds with overwhelming probability.

To carry out the proof, we show that a successful key-revocation adversary allows us to break the uncertainty principle for Gaussian superpositions in Theorem 7.7. The idea is the following: Suppose there exists an algorithm $\mathcal{D}$ that can distinguish Dual-Regev from uniform given ρ_{AUX} in the key-revocation experiment (conditioned on Revoke succeeding on certificate $\mathbf{w}$). Then we can use it to extract (using Goldreich-Levin) a short pre-image $\mathbf{x}_0$ from the internal state ρ_{AUX} which we can send to the challenger as part of the experiment in Theorem 7.7. Moreover, because the revocation certificate $\mathbf{w}$ passes verification[17], it must yield $\mathbf{v}$ when $\mathbf{w}$ is decrypted using a short trapdoor basis (i.e., by running Revoke). Fortunately, once we present a pre-image $\mathbf{x}_0$ in the second half of the experiment, we are allowed to be computationally unbounded; meaning, we can simply decode $\mathbf{w}$ in exponential time (i.e. by breaking LWE, or by finding a trapdoor that can be used to decode $\mathbf{w}$ using Revoke). Therefore, we obtain an adversary that ends up finding both a pre-image $\mathbf{x}_0$ and $\mathbf{v}$ at the same time, which violates Theorem 7.7.

Our first result concerns $(\mathsf{negl}(\lambda),\mathsf{negl}(\lambda))$-security, i.e., where we assume that revocation succeeds with overwhelming probability.

Theorem 7.13. *Let $n\in\mathbb{N}$ and q be a prime with $q=2^{o(n)}$ and $m=\lceil 6n\log q\rceil$, each parameterized by $\lambda\in\mathbb{N}$. Let $q/\sqrt{8m}>\sigma>\sqrt{8m}$ and let $\alpha,\beta\in(0,1)$ be noise ratios chosen such that $\beta/\alpha=2^{o(n)}$ and $1/\alpha=2^{o(n)}\cdot\sigma$. Let $\nu=64m^2$. Then, assuming the subexponential hardness of the $\mathsf{LWE}^m_{n,q,\alpha q}$ and $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ problems, the scheme $\mathsf{CRevDual}=(\mathsf{KeyGen},\mathsf{Enc},\mathsf{Dec},\mathsf{Delete},\mathsf{Revoke})$ in Construction 2 is a $(\mathsf{negl}(\lambda),\mathsf{negl}(\lambda))$-secure key-revocable public-key encryption scheme with classical revocation.*

[17] Here, one can think of it as essentially being an LWE encryption $\mathbf{w}^\intercal=\hat{\mathbf{s}}^\intercal\mathbf{A}+\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}^\intercal+\hat{\mathbf{e}}^\intercal$ of the string $\mathbf{v}$.

Proof. Let $\lambda \in \mathbb{N}$ and suppose that there exists an adversary consisting of a pair of QPT algorithms $(\mathcal{A}, \mathcal{D})$ and a function $\varepsilon(\lambda) = 1/\mathrm{poly}(\lambda)$ such that

$$\left| \Pr\left[\begin{array}{c} \mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top \\ \wedge \\ \mathcal{D}(\mathbf{A},\mathbf{y},\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal,\mathbf{s}^\intercal\mathbf{y}+e',\rho_{\mathrm{AUX}})=1 \end{array} : \begin{array}{c} (\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q) \\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\mathbf{s}\xleftarrow{\$}\mathbb{Z}_q^n \\ \mathbf{e}\sim D_{\mathbb{Z}^m,\,\alpha q},e'\sim D_{\mathbb{Z},\,\beta q} \\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma) \\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle \\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda) \end{array}\right] - \Pr\left[\begin{array}{c} \mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w})=\top \\ \wedge \\ \mathcal{D}(\mathbf{A},\mathbf{y},\mathbf{u},r,\rho_{\mathrm{AUX}})=1 \end{array} : \begin{array}{c} (\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q) \\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\mathbf{u}\xleftarrow{\$}\mathbb{Z}_q^m,r\xleftarrow{\$}\mathbb{Z}_q \\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma) \\ |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle \\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda) \end{array}\right] \right| = \varepsilon(\lambda).$$

We now show how to use $(\mathcal{A}, \mathcal{D})$ to violate the uncertainty principle for Gaussian states from Theorem 7.7. Consider the adversary $(\mathcal{B}_0, \mathcal{B}_1)$ consisting of the following procedures:

- $\mathcal{B}_0(\mathbf{A}, \mathbf{y}, |\psi_{\mathbf{y}}^{\mathbf{v}}\rangle)$: run the algorithm $\mathcal{A}$ to generate

  $$|\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\mathrm{AUX}} \leftarrow \mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}| \otimes \tau_\lambda).$$

  Then, run the Goldreich-Levin extractor $\mathcal{E}(\mathbf{A}, \mathbf{y}, \rho_{\mathrm{AUX}})$ from Theorem 7.12 (instantiated using $\mathcal{A}, \mathcal{D}$) and let $\mathbf{x}_0$ be the outcome. Send $\mathbf{x}_0$ to the challenger and initialize $\mathcal{B}_1$ with $(\mathbf{A}, \mathbf{w})$.

- $\mathcal{B}_1(\mathbf{A}, \mathbf{w})$: compute (in exponential time) a $(20n \log q)$-good trapdoor basis $\hat{\mathbf{T}}_{\mathbf{A}} \in \mathbb{Z}^{m\times m}$ for the matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ as in Definition 7.3 with the properties that

  $$\mathbf{A} \cdot \hat{\mathbf{T}}_{\mathbf{A}} = \mathbf{0} \pmod q \quad \text{and} \quad \|\hat{\mathbf{T}}_{\mathbf{A}}\| \leq 20n \log q.$$

  Then, use $\hat{\mathbf{T}}_{\mathbf{A}}$ to do the following:

  1. Compute $\mathbf{c}^\intercal = \mathbf{w}^\intercal \cdot \hat{\mathbf{T}}_{\mathbf{A}} \pmod q$.
  2. Compute $\mathbf{d}^\intercal = \mathbf{c}^\intercal \cdot \hat{\mathbf{T}}_{\mathbf{A}}^{-1}$, where $\hat{\mathbf{T}}_{\mathbf{A}}^{-1}$ is the inverse matrix of $\hat{\mathbf{T}}_{\mathbf{A}}$ over $\mathbb{R}$.
  3. For $i \in [m]$, let $v_i' = 0$, if $d_i \in [-q/\sigma, q/\sigma]$, and let $v_i' = 1$, otherwise.
  4. Output $\mathbf{v}' = (v_1', \ldots, v_m')$.

Because $\mathcal{D}$ succeeds with distinguishing advantage $\varepsilon = 1/\mathrm{poly}(\lambda)$, it follows from Theorem 7.12 the extractor outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\mathrm{poly}(\lambda, m, \sigma, q, 1/\varepsilon)$ such that

$$\Pr\left[\begin{array}{c} \mathsf{Revoke}(\mathbf{T_A},\mathbf{v},\mathbf{w}) \\ \wedge \\ \mathcal{E}(\mathbf{A},\mathbf{y},\mathrm{AUX}) \in \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0},\sigma\sqrt{\frac{m}{2}}) \end{array} : \begin{array}{c} (\mathbf{A},\mathbf{T_A})\leftarrow\mathsf{GenTrap}(1^n,1^m,q) \\ (|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenGauss}(\mathbf{A},\sigma) \\ \mathbf{v}\xleftarrow{\$}\{0,1\}^m,\,|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\leftarrow\mathbf{Z}_q^{\lfloor\frac{q}{\nu}\rfloor\cdot\mathbf{v}}|\psi_{\mathbf{y}}\rangle \\ |\mathbf{w}\rangle\langle\mathbf{w}|\otimes\rho_{\mathrm{AUX}}\leftarrow\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}(|\psi_{\mathbf{y}}^{\mathbf{v}}\rangle\langle\psi_{\mathbf{y}}^{\mathbf{v}}|\otimes\tau_\lambda) \end{array}\right] \geq \mathrm{poly}(\varepsilon, 1/q).$$

Finally, because $\hat{\mathbf{T}}_{\mathbf{A}} \in \mathbb{Z}^{m\times m}$ is a $(20n \log q)$-good trapdoor basis for $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, this implies that the probability that $\mathcal{B}_1(\mathbf{A}, \mathbf{w})$ correctly recovers $\mathbf{v}$ is also at least $1/\mathrm{poly}(\lambda)$. Therefore, $(\mathcal{B}_0, \mathcal{B}_1)$ has non-negligible probability of finding both the correct certificate $\mathbf{v}' = \mathbf{v}$ as well as a short pre-image $\mathbf{x}_0$ such that $\mathbf{A} \cdot \mathbf{x}_0 = \mathbf{y} \pmod q$ and $\|\mathbf{x}_0\| \leq \sigma\sqrt{m/2}$. This proves the claim. □

Finally, we obtain the following stronger notion of $(\mathsf{negl}(\lambda), 1 - 1/\mathrm{poly}(\lambda))$-security, assuming Theorem 7.9 is true.

Theorem 7.14. *Let $n \in \mathbb{N}$ and q be a prime with $q = 2^{o(n)}$ and $m = \lceil 6n \log q \rceil$, each parameterized by $\lambda \in \mathbb{N}$. Let $q/\sqrt{8m} > \sigma > \sqrt{8m}$ and let $\alpha, \beta \in (0,1)$ be noise ratios chosen such that $\beta/\alpha = 2^{o(n)}$ and $1/\alpha = 2^{o(n)} \cdot \sigma$. Let $\nu = 64m^2$. Then, assuming Theorem 7.9, the scheme $\mathsf{CRevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Delete}, \mathsf{Revoke})$ in Construction 2 is a $(\mathsf{negl}(\lambda), 1-1/\mathrm{poly}(\lambda))$-secure key-revocable public-key encryption scheme with classical revocation.*

Proof. The proof is analogous to Theorem 7.13, except that we invoke Theorem 7.9 instead of Theorem 7.12. □

8 Key-Revocable Fully Homomorphic Encryption

In this section, we describe our key-revocable (leveled) fully homomorphic encryption scheme from LWE which is based on the so-called DualGSW scheme used by Mahadev [Mah18] which itself is a variant of the homomorphic encryption scheme by Gentry, Sahai and Waters [GSW13].

Let $\lambda \in \mathbb{N}$ be the security parameter. Suppose we would like to evaluate L-depth circuits consisting of NAND gates. We choose $n(\lambda, L) \gg L$ and a prime $q = 2^{o(n)}$. Then, for integer parameters $m \geq 2n \log q$ and $N = (m+1) \cdot \lceil \log q \rceil$, we let $\mathbf{I}$ be the $(m+1) \times (m+1)$ identity matrix and let $\mathbf{G} = [\mathbf{I} \,\|\, 2\mathbf{I} \,\|\, \ldots \,\|\, 2^{\lceil \log q \rceil - 1}\mathbf{I}] \in \mathbb{Z}_q^{(m+1)\times N}$ denote the so-called *gadget matrix* which converts a binary representation of a vector back to its original vector representation over the field $\mathbb{Z}_q$. Note that the associated (non-linear) inverse operation $\mathbf{G}^{-1}$ converts vectors in $\mathbb{Z}_q^{m+1}$ to their binary representation in $\{0,1\}^N$. In other words, we have that $\mathbf{G} \circ \mathbf{G}^{-1}$ acts as the identity operator.

8.1 Construction

We now construct a key-revocable Dual-Regev (leveled) fully homomorphic encryption scheme based on our key-revocable Dual-Regev public-key scheme.

Remark 8.1. *While the construction in this section can be readily adapted to feature* classical revocation *via Construction 2, we choose to focus on* quantum revocation *for simplicity by making use of our Dual-Regev scheme from Construction 1.*

We are now ready to state our construction.

Construction 3 (Key-Revocable DualGSW encryption)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. The scheme $\mathsf{RevDualGSW} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Revoke})$ consists of the following* QPT *algorithms:*

$\mathsf{KeyGen}(1^\lambda, 1^L) \to (\mathsf{PK}, \rho_{\mathsf{SK}})$: *sample a pair $(\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ and generate a Gaussian superposition $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ with*

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y}}} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \,,$$

for some $\mathbf{y} \in \mathbb{Z}_q^n$. Output $\mathsf{PK} = (\mathbf{A}, \mathbf{y})$, $\rho_{\mathsf{SK}} = |\psi_{\mathbf{y}}\rangle$ and $\mathsf{MSK} = \mathsf{td}_{\mathbf{A}}$.

$\mathsf{Enc}(\mathsf{PK}, \mu)$: *to encrypt* $\mu \in \{0,1\}$*, parse* $(\mathbf{A}, \mathbf{y}) \leftarrow \mathsf{PK}$*, sample a random matrix* $\mathbf{S} \xleftarrow{\$} \mathbb{Z}_q^{n \times N}$ *and* $\mathbf{E} \sim D_{\mathbb{Z}^{m \times N}, \alpha q}$ *and row vector* $\mathbf{e} \sim D_{\mathbb{Z}^N, \beta q}$*, and output the ciphertext*

$$\mathsf{CT} = \left[\frac{\mathbf{A}^\intercal \mathbf{S} + \mathbf{E}}{\mathbf{y}^\intercal \mathbf{S} + \mathbf{e}}\right] + \mu \cdot \mathbf{G} \pmod q \in \mathbb{Z}_q^{(m+1) \times N}.$$

$\mathsf{Eval}(\mathsf{CT}_0, \mathsf{CT}_1)$: *to apply a* NAND *gate on a ciphertext pair* CT_0 *and* CT_1*, output the matrix*

$$\mathbf{G} - \mathsf{CT}_0 \cdot \mathbf{G}^{-1}(\mathsf{CT}_1) \pmod q \in \mathbb{Z}_q^{(m+1) \times N}.$$

$\mathsf{Dec}(\rho_{\mathsf{SK}}, \mathsf{CT}) \to \{0,1\}$: *to decrypt* CT*, apply the unitary* $U : |\mathbf{x}\rangle |0\rangle \to |\mathbf{x}\rangle |(-\mathbf{x}, 1) \cdot \mathsf{CT}_N\rangle$ *on input* $|\psi_{\mathbf{y}}\rangle \leftarrow \rho_{\mathsf{SK}}$*, where* $\mathsf{CT}_N \in \mathbb{Z}_q^{m+1}$ *is the* N*-th column of* CT*, and measure the second register in the computational basis. Output* 0*, if the measurement outcome is closer to* 0 *than to* $\lfloor \frac{q}{2} \rfloor$*, and output* 1*, otherwise.*

$\mathsf{Revoke}(\mathsf{MSK}, \mathsf{PK}, \rho) \to \{\top, \bot\}$*: on input* $\mathsf{td}_{\mathbf{A}} \leftarrow \mathsf{MSK}$ *and* $(\mathbf{A}, \mathbf{y}) \leftarrow \mathsf{PK}$*, apply the projective measurement* $\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, I - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\}$ *onto* ρ *using* $\mathsf{SampGauss}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ *in Algorithm 2. Output* $\top$ *if the measurement is successful, and output* $\bot$ *otherwise.*

8.2 Proof of security

Our first result on the security of Construction 3 concerns $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda))$-security, i.e., we assume that revocation succeeds with overwhelming probability.

Theorem 8.2. *Let* L *be an upper bound on the* NAND*-depth of the circuit which is to be evaluated. Let* $n \in \mathbb{N}$ *and* q *be a prime modulus with* $n = n(\lambda, L) \gg L$*,* $q = 2^{o(n)}$ *and* $m \geq 2n \log q$*, each parameterized by the security parameter* $\lambda \in \mathbb{N}$*. Let* $N = (m+1) \cdot \lceil \log q \rceil$ *be an integer. Let* $q/\sqrt{8m} > \sigma > \sqrt{8m}$ *and let* $\alpha, \beta \in (0,1)$ *be parameters such that* $\beta/\alpha = 2^{o(n)}$ *and* $1/\alpha = 2^{o(n)} \cdot \sigma$*. Then, assuming the subexponential hardness of the* $\mathsf{LWE}^m_{n,q,\alpha q}$ *and* $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ *problems, the scheme* $\mathsf{RevDualGSW} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Revoke})$ *in Construction 3 is a* $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda))$*-secure key-revocable (leveled) fully homomorphic encryption scheme according to Definition 5.3.*

Proof. Let $\mathcal{A}$ be a QPT adversary and suppose that

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}_{\mathcal{A}}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}_{\mathcal{A}}(1^\lambda, 1)\right]\right| = \epsilon(\lambda),$$

for some $\varepsilon(\lambda)$ with respect to $\mathsf{Expt}_{\mathcal{A}}(1^\lambda, b)$ in Figure 20. Note that the RevDualGSW ciphertext can (up to an additive shift) be thought of as a column-wise concatenation of N-many independent ciphertexts of our key-revocable Dual-Regev scheme in Construction 1. Therefore, we can invoke Theorem 6.1 in order to argue that $\varepsilon(\lambda)$ is at most negligible.

□

Our second result concerns $(\mathsf{negl}(\lambda), 1 - 1/\mathrm{poly}(\lambda))$-security, i.e., we do not make any requirements on the success probability of revocation. Here, we need to invoke Theorem 6.16.

Theorem 8.3. *Let* L *be an upper bound on the* NAND*-depth of the circuit which is to be evaluated. Let* $n \in \mathbb{N}$ *and* q *be a prime modulus with* $n = n(\lambda, L) \gg L$*,* $q = 2^{o(n)}$ *and* $m \geq 2n \log q$*, each parameterized by the security parameter* $\lambda \in \mathbb{N}$*. Let* $N = (m+1) \cdot \lceil \log q \rceil$ *be an integer. Let* $q/\sqrt{8m} >$

$\underline{\mathsf{Expt}_{\mathcal{A}}(1^\lambda, b)}$:

1. The challenger samples $(\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ and generates

$$|\psi_{\mathbf{y}}\rangle \;=\; \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x})\;|\mathbf{x}\rangle\,,$$

   for some $\mathbf{y} \in \mathbb{Z}_q^n$, by running $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$. The challenger lets $\mathsf{MSK} \leftarrow \mathsf{td}_{\mathbf{A}}$ and $\mathsf{PK} \leftarrow (\mathbf{A}, \mathbf{y})$ and sends $\rho_{\mathsf{SK}} \leftarrow |\psi_{\mathbf{y}}\rangle$ to the adversary $\mathcal{A}$.

2. $\mathcal{A}$ generates a (possibly entangled) bipartite state $\rho_{R,\mathrm{AUX}}$ in systems $\mathcal{H}_R \otimes \mathcal{H}_{\mathrm{AUX}}$ with $\mathcal{H}_R = \mathcal{H}_q^m$, returns system R and holds onto the auxiliary system AUX.

3. The challenger runs $\mathsf{Revoke}(\mathsf{MSK}, \mathsf{PK}, \rho_R)$, where ρ_R is the reduced state in system R. If the outcome is $\top$, the game continues. Otherwise, output $\mathsf{Invalid}$.

4. $\mathcal{A}$ submits a plaintext bit $\mu \in \{0, 1\}$.

5. The challenger does the following depending on $b \in \{0, 1\}$:

   - if $b = 0$: The challenger samples a random matrix $\mathbf{S} \xleftarrow{\$} \mathbb{Z}_q^{n\times N}$ and errors $\mathbf{E} \sim D_{\mathbb{Z}^{m\times N},\,\alpha q}$ and row vector $\mathbf{e} \sim D_{\mathbb{Z}^N,\,\beta q}$, and outputs the ciphertext

$$\mathsf{CT} = \left[\begin{array}{c} \mathbf{A}^\intercal\mathbf{S}+\mathbf{E} \\ \hline \mathbf{y}^\intercal\mathbf{S}+\mathbf{e} \end{array}\right] + \mu \cdot \mathbf{G} \;\in \mathbb{Z}_q^{(m+1)\times N}.$$

   - if $b = 1$: the challenger samples a matrix $\mathbf{U} \xleftarrow{\$} \mathbb{Z}_q^{m\times N}$ and row vector $r \xleftarrow{\$} \mathbb{Z}_q^N$ uniformly at random, and sends the following to $\mathcal{A}$:

$$\left[\begin{array}{c} \mathbf{U} \\ \hline \mathbf{r} \end{array}\right] \;\in \mathbb{Z}_q^{(m+1)\times N}.$$

6. $\mathcal{A}$ returns a bit $b' \in \{0, 1\}$.

Figure 20: The key-revocable security experiment according to Definition 5.3.

$\sigma > \sqrt{8m}$ and let $\alpha, \beta \in (0, 1)$ be parameters such that $\beta/\alpha = 2^{o(n)}$ and $1/\alpha = 2^{o(n)} \cdot \sigma$. Then, assuming assuming Theorem 6.16, the scheme $\mathsf{RevDualGSW} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Revoke})$ in Construction 3 is a $(\mathsf{negl}(\lambda), 1 - 1/\mathrm{poly}(\lambda))$-secure key-revocable (leveled) fully homomorphic encryption scheme according to Definition 5.3.

Proof. The proof is the same as in the theorem before, except that we invoke Theorem 6.2 instead of Theorem 6.1 in order to argue security. □

9 Revocable Pseudorandom Functions

In this section, we introduce the notion of *key-revocable* pseudorandom functions (or simply, called *revocable*) and present the first construction from (quantum hardness of) learning with errors.

9.1 Definition

Let us first recall the traditional notion of PRF security [GGM86], defined as follows.

Definition 9.1 (Pseudorandom Function)**.** *Let $\lambda \in \mathbb{N}$ and $\kappa(\lambda), \ell(\lambda)$ and $\ell'(\lambda)$ be polynomials. A (post-quantum) pseudorandom function* (pqPRF) *is a pair* $(\mathsf{Gen}, \mathsf{PRF})$ *of* PPT *algorithms given by*

- $\mathsf{Gen}(1^\lambda)$: *On input 1^λ, it outputs a key $k \in \{0,1\}^\kappa$.*
- $\mathsf{PRF}(k, x)$: *On input $k \in \{0,1\}^\kappa$ and $x \in \{0,1\}^\ell$, it outputs a value $y \in \{0,1\}^{\ell'}$.*

with the property that, for any QPT *distinguisher $\mathcal{D}$, we have*

$$\left|\Pr\left[\mathcal{D}^{\mathsf{PRF}(k,\cdot)}(1^\lambda) = 1\right] : k \leftarrow \mathsf{Gen}(1^\lambda)\right] - \Pr\left[\mathcal{D}^{F(\cdot)}(1^\lambda) = 1\right] : F \xleftarrow{\$} \mathcal{F}^{\ell,\ell'}\right]\Bigg| \leq \mathsf{negl}(\lambda),$$

where $\mathcal{F}^{\ell,\ell'}$ is the set of all functions with domain $\{0,1\}^\ell$ and range $\{0,1\}^{\ell'}$.

We now present a formal definition of revocable pseudorandom functions below.

Definition 9.2 (Revocable Pseudorandom Function)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter and let $\kappa(\lambda), \ell(\lambda)$ and $\ell'(\lambda)$ be polynomials. A revocable pseudorandom function* (rPRF) *is a scheme* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *consisting of the following efficient algorithms:*

- $\mathsf{Gen}(1^\lambda)$*: on input the security parameter $\lambda \in \mathbb{N}$, it outputs a* PRF *key $k \in \{0,1\}^\kappa$, a quantum state ρ_k and a master secret key* MSK*.*
- $\mathsf{PRF}(k, x)$*: on input a key $k \in \{0,1\}^\kappa$ and an input string $x \in \{0,1\}^\ell$, it outputs a value $y \in \{0,1\}^{\ell'}$. This is a deterministic algorithm.*
- $\mathsf{Eval}(\rho_k, x)$*: on input a state ρ_k and an input $x \in \{0,1\}^\ell$, it outputs a value $y \in \{0,1\}^{\ell'}$.*
- $\mathsf{Revoke}(\mathsf{MSK}, \sigma)$*: on input key* MSK *and a state σ, it outputs* Valid *or* Invalid*.*

We additionally require that the following holds:

Correctness. For each $(k, \rho_k, \mathsf{MSK})$ in the support of $\mathsf{Gen}(1^\lambda)$ and for every $x \in \{0,1\}^\ell$:

- (Correctness of evaluation:)

$$\Pr\left[\mathsf{PRF}(k, x) = \mathsf{Eval}(\rho_k, x)\right] \geq 1 - \mathsf{negl}(\lambda).$$

- (Correctness of revocation:)

$$\Pr\left[\mathsf{Valid} \leftarrow \mathsf{Revoke}(\mathsf{MSK}, \rho_k)\right] \geq 1 - \mathsf{negl}(\lambda).$$

$\underline{\mathsf{Expt}_{\mathcal{A},\mu}(1^\lambda, b)}$:

Initialization Phase:

- The challenger computes $(k, \rho_k, \mathsf{MSK}) \leftarrow \mathsf{Gen}(1^\lambda)$ and sends ρ_k to $\mathcal{A}$.

Revocation Phase:

- The challenger sends the message `REVOKE` to $\mathcal{A}$.
- The adversary $\mathcal{A}$ sends a state σ to the challenger.
- The challenger aborts if $\mathsf{Revoke}\,(\mathsf{MSK}, \sigma)$ outputs $\mathsf{Invalid}$.

Guessing Phase:

- The challenger samples bit $b \leftarrow \{0,1\}$.
- The challenger samples random inputs $x_1, \ldots, x_\mu \xleftarrow{\$} \{0,1\}^\ell$ and then sends the values $(x_1, \ldots, x_\mu)$ and $(y_1, \ldots, y_\mu)$ to $\mathcal{A}$, where:
  - If $b = 0$, set $y_1 = \mathsf{PRF}(k, x_1)$, $\ldots$, $y_\mu = \mathsf{PRF}(k, x_\mu)$ and,
  - If $b = 1$, set $y_1, \ldots, y_\mu \xleftarrow{\$} \{0,1\}^{\ell'}$.
- $\mathcal{A}$ outputs a bit b' and wins if $b' = b$.

Figure 21: Revocable PRF security

9.2 Security

We define revocable PRF security below.

Definition 9.3 (Revocable PRF Security)**.** *A revocable pseudorandom function* (rPRF) *consisting of a tuple of* QPT *algorithms* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *has* $(\varepsilon, \delta, \mu)$ *revocable* PRF *security if, for every* QPT *adversary* $\mathcal{A}$ *with*

$$\mathsf{Pr}[\mathsf{Invalid} \leftarrow \mathsf{Expt}_{\mathcal{A},\mu}(1^\lambda, b)] \leq \delta(\lambda)$$

for $b \in \{0,1\}$*, it holds that*

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}_{\mathcal{A},\mu}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}_{\mathcal{A},\mu}(1^\lambda, 1)\right]\right| \leq \varepsilon(\lambda),$$

where $\mathsf{Expt}_{\mathcal{A},\mu}$ *is as defined in Figure 21. If* $\delta(\lambda) = 1 - 1/\mathrm{poly}(\lambda)$*,* $\varepsilon(\lambda) = \mathsf{negl}(\lambda)$ *we oftentimes drop* (δ, ε) *and simply refer to it as* rPRF *satisfies* μ*-revocable* PRF *security.*

From one-query to multi-query security. We show that proving security with respect to $\mu = 1$ is sufficient. That is, we show the following.

Claim 9.4. *Supoose an* rPRF *scheme* (Gen, PRF, Eval, Revoke) *satisfies* 1*-revocable* PRF *security. Then,* rPRF *also satisfies the stronger notion of (multi-query) revocable PRF security.*

Proof. Let $\mathcal{A}$ be a QPT adversary that participating in the revocable PRF security experiment defined in Figure 21 and let $(x_1, y_1), \ldots, (x_\mu, y_\mu)$ denote the challenge input-output pairs, for some polynomial $\mu = \mu(\lambda)$. In the following, we denote by k the PRF key sampled using Gen by the challenger in Figure 21. We consider a sequence of hybrids defined as follows.

H_i, for $i \in [\mu + 1]$: In this hybrid, $y_1, \ldots, y_{i-1}$ are sampled uniformly at random from $\{0,1\}^{\ell'}$ and $y_i, \ldots, y_\mu$ are generated as follows: $y_j = \mathsf{PRF}(k, x_j)$ for $j \geq i$.

We claim that $\mathcal{A}$ cannot distinguish between hybrids H_i and H_{i+1}, for all $i \in [\mu]$, with more than negligible advantage. Suppose for the sake of contradiction that the claim is not true, and that $\mathcal{A}$ can distinguish H_i and H_{i+1}, for some index $i \in [\mu]$, with advantage at least $\varepsilon(\lambda) = 1/\mathrm{poly}(\lambda)$. We will now show that we can use the adversary $\mathcal{A}$ to break the 1-revocation security of rPRF.

Consider a reduction $\mathcal{B}$ that does the following:

1. Receive the state ρ_k from the challenger.
2. Sample $x_{i+1}, \ldots, x_\mu$ uniformly at random from $\{0,1\}^\ell$. Denote $\rho_k^{(i+1)} = \rho_k$. Do the following for $j = i+1, \ldots, \mu$: $\mathsf{Eval}(\rho_k^{(j)}, x_j)$ to obtain y_j. Using the "Almost As Good As New" [Aar16], recover $\rho_k^{(j+1)}$, where $\rho_k^{(j+1)}$ is negligibly[18] close to ρ_k in trace distance.
3. Forward the state $\rho_k^{(\mu+1)}$ to $\mathcal{A}$.
4. When the challenger submits the message `REVOKE`, forward the same message to $\mathcal{A}$.
5. If $\mathcal{A}$ sends σ, then forward the same state σ to the challenger.
6. If the revocation did not fail, the guessing phase begins. The challenger sends (x^*, y^*). Then, sample $x_1, \ldots, x_{i-1}$ uniformly at random from $\{0,1\}^\ell$ and $y_1, \ldots, y_{i-1}$ uniformly at random from $\{0,1\}^{\ell'}$. Set $x_i = x^*$ and $y_i = y^*$. Send $(x_1, y_1), \ldots, (x_\mu, y_\mu)$ to $\mathcal{A}$.
7. Output b, where b is the output of $\mathcal{A}$.

From the quantum union bound (Lemma 2.3), the "Almost As Good As New" lemma (Lemma 2.2) and the correctness of rPRF, it follows that $\mathrm{TD}(\rho_k, \rho_k^{(\mu+1)}) \leq \mathsf{negl}(\lambda)$ and thus, the advantage of $\mathcal{A}$ when given $\rho_k^{(\mu+1)}$ instead of ρ_k is now at least $\varepsilon - \mathsf{negl}(\lambda)$. Moreover, by the design of $\mathcal{B}$, it follows that the success probability of $\mathcal{B}$ in breaking 1-revocation security of rPRF is exactly the same as the success probability of $\mathcal{A}$ in breaking revocation security of rPRF. This contradicts the fact that rPRF satisfies 1-revocation security.

□

Remark 9.5. *Our notion of revocable* PRF *security from Definition 9.3 does not directly imply traditional notion of* pqPRF *security*[19] *from Definition 9.1. The reason is that the definition does*

[18] Technically, this depends on the correctness error and we start with a rPRF that is correct with probability negligibly close to 1.

[19] Although any revocable PRF is a *weak* PRF. Recall that a weak PRF is one where the adversary receives as input $(x_1, y_1), \ldots, (x_\mu, y_\mu)$, where x_is are picked uniformly at random. The goal of the adversary is to distinguish the two cases: all y_is are pseudorandom or all y_is are picked uniformly at random.

not preclude the possibility of there being an input x (say an all zeroes string) on which, PRF *outputs x itself (or the first bit of x if the output of* PRF *is a single bit).*

Motivated by Remark 9.5, we now introduce the following notion of a *strong* rPRF.

Definition 9.6 (Strong rPRF)**.** *We say that a scheme* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *is a strong revocable pseudorandom function (or, strong* rPRF*) if the following two properties hold:*

1. $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *satisfy revocable* PRF *security according to Definition 9.3, and*
2. $(\mathsf{Gen}, \mathsf{PRF})$ *satisfy* pqPRF *security according to Definition 9.1.*

Remark 9.7. *When instantiating pseudorandom functions in the textbook construction of private-key encryption [Gol06] from revocable pseudorandom functions, we immediately obtain a revocable private-key encryption scheme.*

We show that the issue raised in Remark 9.5 is not inherent. In fact, we give a simple generic transformation that allows us to obtain strong rPRFs by making use of traditional pqPRFs.

Claim 9.8 (Generic Transformation for Strong rPRFs)**.** *Let* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *be an* rPRF *scheme which satisfies revocable* PRF *security, and let* $(\overline{\mathsf{Gen}}, \overline{\mathsf{PRF}})$ *be a* pqPRF*. Then, the scheme* $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}}, \widetilde{\mathsf{Eval}}, \widetilde{\mathsf{Revoke}})$ *is a strong* rPRF *which consists of the following algorithms:*

- $\widetilde{\mathsf{Gen}}(1^\lambda)$*: on input the security parameter* 1^λ*, first run* $(k, \rho_k, \mathsf{MSK}) \leftarrow \mathsf{Gen}(1^\lambda)$ *and then output* $((K, k), (K, \rho_k), \mathsf{MSK})$*, where* $K \leftarrow \overline{\mathsf{Gen}}(1^\lambda)$ *is a* pqPRF *key.*
- $\widetilde{\mathsf{PRF}}((K, k), x)$*: on input a key* (K, k) *and string* $x \in \{0,1\}^\ell$*, output* $\overline{\mathsf{PRF}}(K, x) \oplus \mathsf{PRF}(k, x)$*.*
- $\widetilde{\mathsf{Eval}}((K, \rho_k), x)$*: on input* (K, ρ_k) *and* $x \in \{0,1\}^\ell$*, output* $\overline{\mathsf{PRF}}(K, x) \oplus \mathsf{Eval}(\rho_k, x)$*.*
- $\widetilde{\mathsf{Revoke}}(\mathsf{MSK}, (K, \sigma))$*: on input a master secret key* MSK *and a pair* (K, ρ_k)*, first discard the key K and then run* $\mathsf{Revoke}(\mathsf{MSK}, \sigma)$*.*

Proof. Let us first show that the scheme $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}}, \widetilde{\mathsf{Eval}}, \widetilde{\mathsf{Revoke}})$ maintains revocable PRF security. Suppose that there exists a QPT adversary $\mathcal{A}$ and a polynomial $\mu = \mu(\lambda) \in \mathbb{N}$ such that

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}_{\mathcal{A},\mu}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}_{\mathcal{A},\mu}(1^\lambda, 1)\right]\right| = \epsilon(\lambda),$$

for some function $\epsilon(\lambda) = 1/\mathrm{poly}(\lambda)$, and where $\mathsf{Expt}_{\mathcal{A},\mu}$ is the experiment from Figure 21 with respect to the scheme $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}}, \widetilde{\mathsf{Eval}}, \widetilde{\mathsf{Revoke}})$. We show that this implies the existence of a QPT distinguisher $\mathcal{D}$ that breaks the revocable PRF security of the scheme $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$.

The distinguisher $\mathcal{D}$ proceeds as follows:

1. $\mathcal{D}$ receives as input a quantum state ρ_k, where $(k, \rho_k, \mathsf{MSK}) \leftarrow \mathsf{Gen}(1^\lambda)$ is generated by the challenger. Then, $\mathcal{D}$ generates a pqPRF key $K \leftarrow \overline{\mathsf{Gen}}(1^\lambda)$ and sends (K, ρ_k) to $\mathcal{A}$.
2. When $\mathcal{A}$ returns a state ρ, $\mathcal{D}$ forwards it to the challenger as part of the revocation phase.
3. When $\mathcal{D}$ receives the challenge input $(x_1, \ldots, x_\mu)$ and $(y_1, \ldots, y_\mu)$ from the challenger, $\mathcal{D}$ sends $(x_1, \ldots, x_\mu)$ and $(\overline{\mathsf{PRF}}(K, x_1) \oplus y_1, \ldots, \overline{\mathsf{PRF}}(K, x_\mu) \oplus y_\mu)$ to $\mathcal{A}$.

4. When $\mathcal{A}$ outputs b', so does the distinguisher $\mathcal{D}$.

Note that the simulated challenge distribution above precisely matches the challenge distribution from the experiment $\mathsf{Expt}_{\mathcal{A},\mu}$ from Figure 21. Therefore, if $\mathcal{A}$ succeeds with inverse polynomial advantage $\epsilon(\lambda) = 1/\mathrm{poly}(\lambda)$, so does $\mathcal{D}$ – thereby breaking the revocable PRF security of the scheme $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$. Consequently, $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}}, \widetilde{\mathsf{Eval}}, \widetilde{\mathsf{Revoke}})$ satisfies revocable PRF security.

To see why $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}})$ satisfy pqPRF security according to Definition 9.1, we can follow a similar argument as above to break the pqPRF security of $(\overline{\mathsf{Gen}}, \overline{\mathsf{PRF}})$. Here, we rely on the fact that the keys $(k, \rho_k, \mathsf{MSK}) \leftarrow \mathsf{Gen}(1^\lambda)$ and $K \leftarrow \overline{\mathsf{Gen}}(1^\lambda)$ are sampled independently from another.

□

9.3 Construction

We construct a PRF satisfying 1-revocation security (Definition 9.3).

Shift-Hiding Construction. We construct a *shift-hiding* function which is loosely inspired by shift-hiding shiftable functions introduced by Peikert and Shiehian [PS18].

Let $n, m \in \mathbb{N}$, $q \in \mathbb{N}$ be a modulus and let $\ell = nm\lceil \log q \rceil$. In the following, we consider matrix-valued functions $F : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$, where F is one of the following functions:

- $\mathcal{Z} : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$ which, on input $x \in \{0,1\}^\ell$, outputs an all zeroes matrix $\mathbf{0} \in \mathbb{Z}_q^{n \times m}$, or:
- $H_r : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$ which, on input $x \in \{0,1\}^\ell$, outputs $\mathbf{M} \in \mathbb{Z}_q^{n \times m}$, where $r \in \{0,1\}^\ell$ and $x = r \oplus \mathsf{bindecomp}(\mathbf{M})$, where $\mathbf{M} \in \mathbb{Z}_q^{n \times m}$ and $\mathsf{bindecomp}(\cdot)$ takes as input a matrix and outputs a binary string that is obtained by concatenating the binary decompositions of all the elements in the matrix (in some order).

We show that there exist PPT algorithms $(\mathcal{KG}, \mathcal{E})$ (formally defined in Construction 4) with the following properties:

- $\mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$: on input $1^n, 1^m$, a modulus $q \in \mathbb{N}$, a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a function $F \in \{\mathcal{Z}\} \cup \{H_r : r \in \{0,1\}^\ell\}$, it outputs a pair of keys (pk_F, sk_F).
- $\mathcal{E}(pk_F, x)$: on input pk_F, $x \in \{0,1\}^\ell$, it outputs $\mathbf{S}_x\mathbf{A} + \mathbf{E}_x + F(x)$, where $\mathbf{S}_x \in \mathbb{Z}_q^{n \times n}$ and $\mathbf{E}_x \in \mathbb{Z}_q^{n \times m}$, where $||\mathbf{E}_x|| \leq nm^2\sigma\lceil \log(q) \rceil$. Moreover, there is an efficient algorithm that recovers $\mathbf{S}_x$ given sk_F and x.

We show that our construction of $(\mathcal{KG}, \mathcal{E})$ satisfies a *shift-hiding property*; namely, for any $r \in \{0,1\}^\ell$,

$$\{pk_{\mathcal{Z}}\} \approx_c \{pk_{H_r}\},$$

for any pk_F with $(pk_F, sk_F) \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$, where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, and $F \in \{\mathcal{Z}, H_r\}$.

In the construction below, we consider a bijective function $\phi : [n] \times [m] \times [\lceil \log(q) \rceil] \to [\ell]$.

Construction 4. *Consider the* PPT *algorithms* $(\mathcal{KG}, \mathcal{E})$ *defined as follows:*

- $\mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$*: on input* $1^n, 1^m$*, a modulus* $q \in \mathbb{N}$*, a matrix* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ *and function* $F \in \{\mathcal{Z}\} \cup \{H_r : r \in \{0,1\}^\ell\}$*, it outputs a pair of keys* $\kappa_F = (pk_F, sk_F)$ *generated as follows:*

1. *For every* $i, j \in [n], \tau \in [\lceil \log(q) \rceil]$, *define* $\{\mathsf{M}_b^{(i,j,\tau)}\}$ *as follows:*
   - *If* $F = \mathcal{Z}$, *then for every* $i \in [n], j \in [m], \tau \in [\lceil \log(q) \rceil]$, *let* $\mathsf{M}_b^{(i,j,\tau)} = \mathbf{0} \in \mathbb{Z}_q^{n \times n}$,
   - *If* $F = H_r$, *then for every* $i \in [n], j \in [m], \tau \in [\lceil \log(q) \rceil]$, *let* $\mathsf{M}_b^{(i,j,\tau)} = (b \oplus r_{\phi(i,j,\tau)}) \cdot \mathbf{I}_{n \times n}$.
2. *For every* $i \in [n], j \in [m], \tau \in [\lceil \log(q) \rceil]$, $b \in \{0,1\}$, *compute:*

$$pk_b^{(i,j,\tau)} = \mathbf{S}_b^{(i,j,\tau)} \mathbf{A} + \mathbf{E}_b^{(i,j,\tau)} + \mathsf{M}_b^{(i,j,\tau)},$$

$$sk_b^{(i,j,\tau)} = \left(\left\{ \mathbf{S}_b^{(i,j,\tau)}, \mathbf{E}_b^{(i,j,\tau)} \right\} \right),$$

   where for every $i \in [n], j \in [m], \tau \in [\lceil \log(q) \rceil]$, $b \in \{0,1\}$*:*
   - $\mathbf{S}_b^{(i,j,\tau)} \leftarrow D_{\mathbb{Z}_q,\sigma}^{n \times n}$,
   - $\mathbf{E}_b^{(i,j,\tau)} \leftarrow D_{\mathbb{Z}_q,\sigma}^{n \times m}$
3. *Output* $pk_F = \left(\mathbf{A}, \left\{ pk_b^{(i,j,\tau)} \right\}_{\substack{i \in [n], j \in [m], \\ \tau \in [\lceil \log(q) \rceil], b \in \{0,1\}}} \right)$ *and* $sk_F = \left\{ sk_b^{(i,j,\tau)} \right\}_{\substack{i \in [n], j \in [m], \\ \tau \in [\lceil \log(q) \rceil], b \in \{0,1\}}}$.

- $\mathcal{E}(pk_F, x)$*: on input* pk_F *and* $x \in \{0,1\}^\ell$*, proceed as follows:*
  1. *Parse* $pk_F = \left(\mathbf{A} \left\{ pk_b^{(i,j,\tau)} \right\}_{\substack{i \in [n], j \in [m], \\ \tau \in [\lceil \log(q) \rceil], b \in \{0,1\}}} \right)$
  2. *Output* $\sum_{\substack{i \in [n], j \in [m], \\ \tau \in [\lceil \log(q) \rceil]}} pk_{x_{\phi(i,j,\tau)}}^{(i,j,\tau)}$.

Claim 9.9 (Correctness)**.** *Let* $(\mathcal{KG}, \mathcal{E})$ *be the pair of* PPT *algorithms in Construction 4. Let* $(pk_F, sk_F) \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$ *with* $F \in \{\mathcal{Z}\} \cup \{H_r : r \in \{0,1\}^\ell\}$*. Then, the output of* $\mathcal{E}(pk_F, x)$ *is of the form:*

$$\mathcal{E}(pk_F, x) = \mathbf{S}_x \mathbf{A} + \mathbf{E}_x + F(x),$$

where $\mathbf{S}_x \in \mathbb{Z}_q^{n \times n}$ *and* $\mathbf{E}_x \in \mathbb{Z}_q^{n \times m}$ *with* $||\mathbf{E}_x|| \leq nm^2 \sigma \lceil \log(q) \rceil$*. Moreover, there is an efficient algorithm that recovers* $\mathbf{S}_x$ *given* (pk_F, sk_F)*.*

Proof. Let $(pk_F, sk_F) \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$. Parse $pk_F = \left(\mathbf{A}, \left\{ pk_b^{(i,j,\tau)} \right\}_{\substack{i \in [n], j \in [m], \\ \tau \in [\lceil \log(q) \rceil], b \in \{0,1\}}} \right)$ and $sk_F = \left\{ sk_b^{(i,j,\tau)} \right\}_{\substack{i \in [n], j \in [m], \\ \tau \in [\lceil \log(q) \rceil], b \in \{0,1\}}}$, where:

$$pk_b^{(i,j,\tau)} = \mathbf{S}_b^{(i,j,\tau)} \mathbf{A} + \mathbf{E}_b^{(i,j,\tau)} + \mathsf{M}_b^{(i,j,\tau)},$$

$$sk_b^{(i,j,\tau)} = \left(\{ \mathbf{S}_b^{(i,j,\tau)}, \mathbf{E}_b^{(i,j,\tau)} \} \right)$$

There are two cases to consider here:

Case 1. $F = \mathcal{Z}$: in this case, $\mathsf{M}_b^{(i,j,\tau)} = \mathbf{0}$, for every $i \in [n], j \in [m], \tau \in [\lceil \log(q) \rceil], b \in \{0,1\}$. Thus, the following holds:

$$\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} pk^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}} = \underbrace{\left(\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} \mathbf{S}^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}}\right)}_{\mathbf{S}_x}\mathbf{A} + \underbrace{\left(\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} \mathbf{E}^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}}\right)}_{\mathbf{E}_x} + \left(\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} \mathsf{M}^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}}\right)$$
$$= \mathbf{S}_x\mathbf{A} + \mathbf{E}_x + \mathcal{Z}(x)$$

Moreover, we have that $||\mathbf{E}_b^{(i,j,\tau)}|| \leq m\sigma$ and thus, $||\mathbf{E}_x|| \leq nm^2\sigma\lceil\log(q)\rceil$.

Case 2. $F = H_r$:

$$\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} pk^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}} = \mathbf{S}_x\mathbf{A} + \mathbf{E}_x + \left(\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} \mathsf{M}^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}}\right)$$
$$= \mathbf{S}_x\mathbf{A} + \mathbf{E}_x + H_r(x),$$

where $\mathbf{S}_x$ and $\mathbf{E}_x$ are as defined above. The second equality holds because of the fact that $\mathsf{M}^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}}$ has the value $(b \oplus r_{\phi(i,j,\tau)}) \cdot 2^\tau$ in the $(i,j)^{th}$ position and zero, everywhere else. Thus, summing up all the $\mathsf{M}^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}}$ matrices results in the matrix M, where $x \oplus r$ is the binary decomposition of M.

Finally, it is clear that $\mathbf{S}_x$ can be efficiently recovered from sk_F and x. □

Claim 9.10 (Shift-hiding property)**.** *Assuming the quantum hardness of learning with errors, the pair $(\mathcal{KG}, \mathcal{E})$ in Construction 4 has the property that*

$$\{pk_{\mathcal{Z}}\} \approx_c \{pk_{H_r}\},$$

for any pk_F with $(pk_F, sk_F) \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$, where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, $r \in \{0,1\}^\ell$ and $F \in \{\mathcal{Z}, H_r\}$.

Proof. For every $i \in [n], j \in [m], \tau \in [\lceil\log(q)\rceil]$, $b \in \{0,1\}$, let $\mathsf{M}_b^{(i,j,\tau)} = (b \oplus r_{\phi(i,j,\tau)}) \cdot \mathbf{I}_{n\times n}$. Then from the quantum hardness of learning with errors, the following holds for every (i,j,τ) and $b \in \{0,1\}$:

$$\{\mathbf{S}_b^{(i,j,\tau)}\mathbf{A} + \mathbf{E}_b^{(i,j,\tau)}\} \approx_c \{\mathbf{S}_b^{(i,j,\tau)}\mathbf{A} + \mathbf{E}_b^{(i,j,\tau)} + \mathsf{M}_b^{(i,j,\tau)}\}$$

Since $\{\mathbf{S}_b^{(i,j,\tau)}\}$ and $\{\mathbf{E}_b^{(i,j,\tau)}\}$ are sampled independently for every (i,j,τ) and $b \in \{0,1\}$, the proof of the claim follows. □

Remark 9.11. *When consider the all-zeroes function $\mathcal{Z}$, we drop the notation from the parameters. For instance, we denote $pk_{\mathcal{Z}}$ to be simply pk.*

Construction. We consider the following parameters which are relevant to our PRF construction. Let $n, m \in \mathbb{N}$ and let $q \in \mathbb{N}$ be a modulus with $q = 2^{o(n)}$, and let $\ell = nm\lceil \log q \rceil$. Let σ be a parameter with $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and let $p \ll q$ be a sufficiently large rounding parameter with

$$n \cdot m^3 \sigma^2 \lceil \log q \rceil = (q/p) \cdot 2^{-o(n)}.$$

We now construct a key-revocable Dual-Regev (leveled) fully homomorphic encryption scheme based on our key-revocable Dual-Regev public-key scheme.

Remark 9.12. *While the construction in this section can be readily adapted to feature* classical revocation *via Construction 2, we choose to focus on* quantum revocation *for simplicity by making use of our Dual-Regev scheme from Construction 1.*

We describe our construction below.

Construction 5 (Revocable PRF scheme)**.** *Let $n \in \mathbb{N}$ be the security parameter and $m \in \mathbb{N}$. Let $q \geq 2$ be a prime and let $\sigma > 0$ be a parameter. Let $(\mathcal{KG}, \mathcal{E})$ be the procedure in Construction 4. Our revocable* PRF *scheme is defined as follows:*

- $\mathsf{Gen}(1^\lambda)$*: This is the following key generation procedure:*
  1. *Sample* $(\mathbf{A}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$.
  2. *Compute* $\kappa_{\mathcal{Z}} \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, \mathcal{Z})$*, where* $\mathcal{Z} : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$ *is the such that* $\mathcal{Z}(x)$ *outputs an all zero matrix for every* $x \in \{0,1\}^\ell$*. Parse* $\kappa_{\mathcal{Z}}$ *as* (pk, sk).
  3. *Generate a Gaussian superposition* $(|\psi_{\mathbf{y}}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ *with*

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{Ax} = \mathbf{y}}} \rho_\sigma(\mathbf{x}) |\mathbf{x}\rangle .$$

  Output $k = (pk, sk, \mathbf{y})$*,* $\rho_k = (pk, |\psi_{\mathbf{y}}\rangle)$ *and* $\mathsf{MSK} = \mathsf{td}_{\mathbf{A}}$.

- $\mathsf{PRF}(k, x)$*: this is the following procedure:*
  1. *Parse the key* k *as a tuple* $(pk, sk), \mathbf{y})$.
  2. *Output* $\lfloor \mathbf{S}_x \mathbf{y} \rceil_p$*. Here,* $\mathbf{S}_x \in \mathbb{Z}_q^{n \times n}$ *is a matrix that can be efficiently recovered from* sk *as stated in Claim 9.9.*

- $\mathsf{Eval}(\rho_k, x)$*: this is the following evaluation algorithm:*
  1. *Parse* ρ_k *as* (pk, ρ).
  2. *Compute* $\mathsf{M}_x \leftarrow \mathcal{E}(pk, x)$.
  3. *Measure the register* Aux *of the state* $U(\rho \otimes |0\rangle\langle 0|_{\mathsf{Aux}})U^\dagger$*. Denote the resulting outcome to be* $\mathbf{z}$*, where* U *is defined as follows:*

$$U |\mathbf{t}\rangle |0\rangle_{\mathsf{Aux}} \to |\mathbf{t}\rangle |\lfloor \mathsf{M}_x \cdot \mathbf{t} \rceil_p\rangle_{\mathsf{Aux}}$$

  4. *Output* $\mathbf{z}$.

- $\mathsf{Revoke}(\mathsf{MSK}, \rho)$: *given as input the trapdoor* $\mathsf{td}_{\mathbf{A}} \leftarrow \mathsf{MSK}$*, apply the projective measurement* $\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, I - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\}$ *onto the state* ρ *using the procedure* $\mathsf{QSampGauss}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ *in Algorithm 2. Output* Valid *if the measurement is successful, and* $\mathsf{Invalid}$ *otherwise.*

Lemma 9.13. *The above scheme satisfies correctness for our choice of parameters.*

Proof. The correctness of revocation follows immediately from the correctness of $\mathsf{QSampGauss}$ in Algorithm 2, which we showed in Theorem 3.3. Next, we show the correctness of evaluation. Let $\kappa_{\mathcal{Z}} \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, \mathcal{Z})$ with $\kappa_{\mathcal{Z}} = (\mathsf{PK}, \mathsf{SK})$. From Claim 9.9, we have for any $x \in \{0,1\}^{\ell}$:

$$\mathcal{E}(\mathsf{PK}, x) = \mathbf{S}_x \mathbf{A} + \mathbf{E}_x \pmod q,$$

where $\mathbf{S}_x \in \mathbb{Z}_q^{n \times n}$ and $\mathbf{E}_x \in \mathbb{Z}_q^{n \times m}$ with $||\mathbf{E}_x||_\infty \leq nm^2\sigma\lceil\log(q)\rceil$. Recall that $\mathsf{GenGauss}(\mathbf{A}, \sigma)$ outputs a state $|\psi_{\mathbf{y}}\rangle$ that is overwhelmingly supported on vectors $\mathbf{t} \in \mathbb{Z}_q^m$ such that $\|\mathbf{t}\| \leq \sigma\sqrt{\frac{m}{2}}$ with $\mathbf{A} \cdot \mathbf{t} = \mathbf{y} \pmod q$. Therefore, we have for any input $x \in \{0,1\}^{\ell}$:

$$\lfloor \mathcal{E}(\mathsf{PK}, x) \cdot \mathbf{t} \rceil_p = \lfloor \mathbf{S}_x \mathbf{A} \cdot \mathbf{t} + \mathbf{E}_x \cdot \mathbf{t} \rceil_p = \lfloor \mathbf{S}_x \cdot \mathbf{y} + \mathbf{E}_x \cdot \mathbf{t} \rceil_p = \lfloor \mathbf{S}_x \cdot \mathbf{y} \rceil_p ,$$

where the last equality follows from the fact that

$$\|\mathbf{E}_x \cdot \mathbf{t}\|_2 \leq \|\mathbf{E}_x\|_2 \cdot \|\mathbf{t}\|_2 \leq \sqrt{m} \cdot \|\mathbf{E}_x\| \cdot \|\mathbf{t}\|_2 \leq n\sqrt{m}m^2\sigma\lceil\log(q)\rceil \cdot \sigma\sqrt{m/2}.$$

and $n \cdot m^3\sigma^2\lceil\log q\rceil = (q/p) \cdot 2^{-o(n)}$ for our choice of parameters. □

Proof of security. Our first result on the security of Construction 5 concerns $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda), 1)$-security, i.e., we assume that revocation succeeds with overwhelming probability.

Theorem 9.14. *Let* $n \in \mathbb{N}$ *and* q *be a prime modulus with* $q = 2^{o(n)}$ *and* $m \geq 2n\log q$*, each parameterized by* $\lambda \in \mathbb{N}$*. Let* $\ell = nm\lceil\log q\rceil$*. Let* $\sqrt{8m} < \sigma < q/\sqrt{8m}$ *and* $\alpha \in (0,1)$ *be any noise ratio with* $1/\alpha = \sigma \cdot 2^{o(n)}$*, and let* $p \ll q$ *be a sufficiently large rounding parameter with*

$$n \cdot m^3\sigma^2\lceil\log q\rceil = (q/p) \cdot 2^{-o(n)}.$$

Then, assuming the quantum subexponential hardness of $\mathsf{LWE}^m_{n,q,\alpha q}$ *and* $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$*, our revocable* PRF *scheme* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *defined in Construction 5 satisfies* $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda), 1)$*-revocation security according to Definition 9.3.*

Proof. Let $\mathcal{A}$ be a QPT adversary and suppose that

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}^{\mathcal{A}}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}^{\mathcal{A}}(1^\lambda, 1)\right]\right| = \epsilon(\lambda),$$

for some function $\varepsilon(\lambda)$ with respect to security experiment $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, b)$ from Figure 22. To complete the proof, it suffices to show that $\varepsilon(\lambda)$ is negligible.

Suppose for the sake of contradition that $\epsilon(\lambda) = 1/\mathrm{poly}(\lambda)$. Let us now introduce a sequence of hybrid experiments which will be relevant for the remainder of the proof.

Let $\mathsf{RevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ be the n-bit key-revocable Dual-Regev scheme from Construction 1. Fix $\mu = 0^n$, where μ is the challenge message in the dual-Regev encryption security.

H_0: This is $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, 0)$ in Figure 22.

H_1: This is the same experiment as $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, 0)$, except for the following changes:

$\underline{\mathsf{Expt}^{\mathcal{A}}(1^\lambda, b)}$:

Initialization Phase:

- The challenger runs the procedure $\mathsf{Gen}(1^\lambda)$:
  1. Sample $(\mathbf{A}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$.
  2. Generate $\mathbf{A}_N \in \mathbb{Z}_q^{(n+m)\times m}$ with $\overline{\mathbf{A}_N} \xleftarrow{\$} \mathbb{Z}_q^{m\times m}$ and $\underline{\mathbf{A}_N} = \mathbf{A}$.
  3. Compute $\kappa_{\mathcal{Z}} \leftarrow \mathcal{KG}(1^n, 1^m, 1^q, \mathbf{A}_N, \mathcal{Z})$, where $\mathcal{KG}$ is as defined in Construction 4 and $\mathcal{Z} : \{0,1\}^\ell \to \mathbb{Z}_q^{n\times m}$ is such that $\mathcal{Z}(x)$ outputs an all zero matrix for every $x \in \{0,1\}^\ell$. Parse $\kappa_{\mathcal{Z}}$ as (pk, sk).
  4. Generate $(|\psi_{\mathbf{y}}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ with
  $$|\psi_{\mathbf{y}}\rangle \;=\; \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle\,.$$
  5. Let $k = (pk, sk, \mathbf{y})$, $\rho_k = (pk, |\psi_{\mathbf{y}}\rangle)$ and $\mathsf{MSK} = \mathsf{td}_{\mathbf{A}}$.
- The challenger sends $\rho_k = (pk, |\psi_{\mathbf{y}}\rangle)$ to $\mathcal{A}$.

Revocation Phase:

- The challenger sends the message `REVOKE` to $\mathcal{A}$.
- $\mathcal{A}$ generates a (possibly entangled) bipartite quantum state $\rho_{R,\text{AUX}}$ in systems $\mathcal{H}_R \otimes \mathcal{H}_{\text{AUX}}$ with $\mathcal{H}_R = \mathcal{H}_q^m$, returns system R and holds onto the auxiliary system AUX.
- The challenger runs $\mathsf{Revoke}(\mathsf{MSK}, \rho_R)$, where ρ_R is the reduced state in system R. If the outcome is $\mathsf{Invalid}$, the challenger aborts.

Guessing Phase:

- The challenger samples $x \leftarrow \{0,1\}^\ell$ and sends (x, y) to $\mathcal{A}$, where
  - If $b = 0$: compute $\mathbf{S}_x$ from sk as in Claim 9.9. Set $y = \lfloor \mathbf{S}_x \mathbf{y} \rceil_p$.
  - If $b = 1$: sample $y \leftarrow \{0,1\}^n$.
- $\mathcal{A}$ outputs a string b' and wins if $b' = b$.

Figure 22: The revocable PRF experiment $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, b)$ for Construction 5.

- Sample a random string $r \leftarrow \{0,1\}^\ell$.
- Run the procedure $\mathsf{RevDual.KeyGen}(1^\lambda)$ instead of $\mathsf{GenTrap}(1^n, 1^m, q)$ and $\mathsf{GenGauss}(\mathbf{A}, \sigma)$ to obtain $(\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathbf{y} \in \mathbb{Z}_q^n, \mathsf{MSK}, \rho_{\mathsf{SK}})$.
- Compute $(\mathsf{CT}_1, \mathsf{CT}_2) \leftarrow \mathsf{RevDual.Enc}(\mathbf{A}, \mathbf{y}, \mu)$, where $\mathsf{CT}_1 \in \mathbb{Z}_q^{n\times m}$ and $\mathsf{CT}_2 \in \mathbb{Z}_q^n$.
- Set $x = r \oplus \mathsf{bindecomp}(\mathsf{CT}_1)$.

The rest of the hybrid is the same as before.

Note that Hybrids H_0 and H_1 are identically distributed.

H_2: This is the same experiment as before, except that the challenger now uses an alternative key-generation algorithm:

- As before, run the procedure $\mathsf{RevDual.KeyGen}(1^\lambda)$ instead of $\mathsf{GenTrap}(1^n, 1^m, q)$ and $\mathsf{GenGauss}(\mathbf{A}, \sigma)$ to obtain $(\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathbf{y} \in \mathbb{Z}_q^n, \mathsf{MSK}, \rho_{\mathsf{SK}})$. Sample $r \leftarrow \{0,1\}^\ell$.
- Let $H_r : \{0,1\}^\ell \to \mathbb{Z}_q^{n\times m}$ be as defined in the beginning of Section 9.3.
- Run the alternate algorithm $\kappa_H \leftarrow \mathcal{KG}(1^n, 1^m, 1^q, \mathbf{A}, H_r)$ instead of $\kappa_{\mathcal{Z}} \leftarrow \mathcal{KG}(1^n, 1^m, 1^q, \mathbf{A}, \mathcal{Z})$.
- Compute the ciphertext $(\mathsf{CT}_1^*, \mathsf{CT}_2^*) \leftarrow \mathsf{RevDual.Enc}(\mathbf{A}, \mathbf{y}, \mu)$, where $\mathsf{CT}_1^* \in \mathbb{Z}_q^{n\times m}$. Then, set $x^* = r \oplus \mathsf{bindecomp}(\mathsf{CT}_1^*)$. Send x^* to the adversary in the guessing phase.

H_3: This is the same hybrid as before, except that we choose $\mathsf{CT}_1^* \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and $\mathsf{CT}_2^* \xleftarrow{\$} \mathbb{Z}_q^n$.

H_4: This is the $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, 1)$ in Figure 22.

Note that hybrids H_3 and H_4 are identically distributed.

We now show the following.

Claim 9.15. *By the shift-hiding property of $(\mathcal{KG}, \mathcal{E})$ in Claim 9.10, we have that the two hybrids H_1 and H_2 are computationally indistinguishable, i.e.*

$$\mathsf{H}_1 \approx_c \mathsf{H}_2.$$

Proof. Suppose for the sake of contradiction that there exists a non-negligble difference in the advantage of the adversary $\mathcal{A}$ in the two hybrids H_1 and H_2.

We now design a reduction $\mathcal{B}$ that violates the shift-hiding property as follows.

1. Sample $r \xleftarrow{\$} \{0,1\}^\ell$. Send $(\mathcal{Z}, H_r)$ to the challenger.
2. The challenger responds with $pk = \left(\mathbf{A}, \left\{\mathsf{CT}_b^{(i,j,\tau)}\right\}_{\substack{i\in[n], j\in[m] \\ \tau\in[\lceil\log(q)\rceil], b\in\{0,1\}}}\right)$
3. Compute $(|\psi_{\mathbf{y}}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ from the challenger.
4. Set $\rho_k = (pk, \rho)$.

5. Compute $(\mathsf{CT}_1, \mathsf{CT}_2) \leftarrow \mathsf{RevDual.Enc}(\mathbf{A}, \mathbf{y}, \mu)$, where $\mathsf{CT}_1 \in \mathbb{Z}_q^{n \times m}$ and $\mathsf{CT}_2 \in \mathbb{Z}_q^n$. Then, set $x^* = r \oplus \mathsf{bindecomp}(\mathsf{CT}_1)$.

6. Compute $\mathsf{Eval}(\rho_k, x^*)$ to obtain y^* while recovering ρ_k^* (using the "Almost as Good As New Lemma", Lemma 2.2) such that $\mathsf{TD}(\rho_k^*, \rho_k) \leq \mathsf{negl}(\lambda)$.

7. Send ρ_k^* to $\mathcal{A}$.

8. $\mathcal{A}$ computes a state on two registers R and AUX. It returns the state on the register R.

9. $\mathcal{A}$, on input the register AUX and (x^*, y^*), outputs a bit b'.

10. Output b'.

If pk is generated using $\mathcal{KG}(1^n, 1^m, q, \mathbf{A}, \mathcal{Z})$ then we are precisely in H_1 (except that Revoke is not performed). Moreover, if pk is generated using $\mathcal{KG}(1^n, 1^m, q, \mathbf{A}, H_r)$ then we are in the hybrid H_2 (except that Revoke is not performed). Therefore $\mathcal{B}$ has a non-negligible advantage at distinguishing H_1 and H_2 whenever Revoke outputs $\top$ on system R. Using Lemma 6.10, this implies that we can break the shift-hiding property with non-negligible advantage. This proves the claim. □

Next, we invoke the security of the n-bit variant of our key-revocable Dual-Regev scheme (which is implied by Theorem 6.1) to show the following.

Claim 9.16. *By the security of our n-bit key-revocable Dual-Regev encryption scheme, we have that the two hybrids H_2 and H_3 are computationally indistinguishable, i.e.*

$$\mathsf{H}_2 \approx_c \mathsf{H}_3.$$

Proof. Suppose for the sake of contradiction that there exists a non-negligble difference in the advantage of $\mathcal{A}$ in the two hybrids H_2 and H_3. Using $\mathcal{A}$, we can now design a reduction $\mathcal{B}$ that violates the revocation security of our n-bit revocable Dual-Regev scheme which is implicit in Theorem 6.1.

The reduction $\mathcal{B}$ proceeds as follows.

1. First, it receives as input $\mathbf{A}, \mathbf{y}$ and a quantum state

$$|\psi_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{Ax}=\mathbf{y}}} \rho_\sigma(\mathbf{x}) |\mathbf{x}\rangle .$$

2. The reduction generates a quantum state ρ_k as follows:

   - Sample a random string $r \xleftarrow{\$} \{0,1\}^\ell$.
   - Let $H_r : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$ be as defined in the beginning of Section 9.3.
   - Run the algorithm $\kappa_H \leftarrow \mathcal{KG}(1^n, 1^m, 1^q, \mathbf{A}, H_r)$ and parse κ_H as (pk, sk).
   - Set $\rho_k = (pk, |\psi_{\mathbf{y}}\rangle)$.

   Send ρ_k to $\mathcal{A}$.

3. $\mathcal{A}$ outputs a state on two registers R and AUX. The register R is returned. The reduction forwards the register R to the challenger.

4. The reduction then gets the challenge ciphertext $\mathsf{CT} = [\mathsf{CT}_1, \mathsf{CT}_2]^\intercal \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$. The reduction then sets

$$x^* = r \oplus \mathsf{bindecomp}(\mathsf{CT}_1)$$

and sends x^* to $\mathcal{A}$ in the guessing phase, together with $y = \lfloor \mathbf{S}_{x^*}\mathbf{y} + \mathsf{CT}_2 \rceil_p$ which is computed using the secret key SK (c.f. Claim 9.9).

5. $\mathcal{A}$ outputs a bit b'. $\mathcal{B}$ outputs b'.

There are two cases to consider here. In the first case, we have $\mathsf{CT} = [\mathsf{CT}_1, \mathsf{CT}_2]^\intercal \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$ is a Dual-Regev ciphertext. Here, $y = \lfloor \mathbf{S}_{x^*}\mathbf{y} + \mathsf{CT}_2 \rceil_p$ precisely corresponds to the output of the pseudorandom function on ρ_k and x. In the second case, we have $\mathsf{CT} = [\mathsf{CT}_1, \mathsf{CT}_2]^\intercal$, where $\mathsf{CT}_1 \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ and $\mathsf{CT}_2 \xleftarrow{\$} \mathbb{Z}_q^m$. Therefore, the resulting string $y = \lfloor \mathbf{S}_{x^*}\mathbf{y} + \mathsf{CT}_2 \rceil_p$ is negligibly close (in total variation distance) to a uniform distribution on $\mathbb{Z}_p^m$.

Putting everything together, we find that the first case corresponds precisely to H_2, whereas the second case corresponds to H_3. As a result, $\mathcal{B}$ violates the revocation security of our n-bit revocable Dual-Regev scheme which is implicit in Theorem 6.1.This completes the proof. □

Putting everything together, we have shown that

$$\left| \Pr\left[1 \leftarrow \mathsf{Expt}^{\mathcal{A}}(1^\lambda, 0) \right] - \Pr\left[1 \leftarrow \mathsf{Expt}^{\mathcal{A}}(1^\lambda, 1) \right] \right| \leq \mathsf{negl}(\lambda).$$

□

Finally, we prove that Construction 5 achieves the stronger notion of $(\mathsf{negl}(\lambda), 1 - 1/\mathrm{poly}(\lambda), 1)$-security assuming Theorem 6.16.

Theorem 9.17. *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and $m \geq 2n \log q$, each parameterized by $\lambda \in \mathbb{N}$. Let $\ell = nm\lceil \log q \rceil$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and $\alpha \in (0, 1)$ be any noise ratio with $1/\alpha = \sigma \cdot 2^{o(n)}$, and let $p \ll q$ be a sufficiently large rounding parameter with*

$$n \cdot m^3 \sigma^2 \lceil \log q \rceil = (q/p) \cdot 2^{-o(n)}.$$

Then, assuming Theorem 6.16, our revocable PRF scheme $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ defined in Construction 5 has $(\mathsf{negl}(\lambda), 1 - 1/\mathrm{poly}(\lambda), 1)$-revocation security according to Definition 9.3.

Proof. The proof is the same as in Theorem 9.14, except that we invoke Theorem 6.2 instead of Theorem 6.1 for Dual-Regev security. □

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 pp. 4, 16).

[Aar16] Scott Aaronson. *The Complexity of Quantum States and Transformations: From Quantum Money to Black Holes*. 2016. arXiv: 1607.05256 [quant-ph] (cit. on pp. 9, 19, 75).

[AC02] Mark Adcock and Richard Cleve. "A quantum Goldreich-Levin theorem with cryptographic applications". In: *STACS 2002*. Ed. by Helmut Alt and Afonso Ferreira. Springer, 2002, pp. 323–334. ISBN: 978-3-540-45841-8. DOI: 10.1007/3-540-45841-7_26. arXiv: quant-ph/0108095 (cit. on p. 11).

[Ajt96] Miklós Ajtai. "Generating Hard Instances of Lattice Problems (Extended Abstract)". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*. Ed. by Gary L. Miller. ACM, 1996, pp. 99–108. DOI: 10.1145/237814.237838. URL: https://doi.org/10.1145/237814.237838 (cit. on pp. 10, 22).

[AK21] Prabhanjan Ananth and Fatih Kaleoglu. "Unclonable encryption, revisited". In: *Theory of Cryptography: 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8–11, 2021, Proceedings, Part I*. Springer. 2021, pp. 299–329 (cit. on p. 17).

[AK22] Prabhanjan Ananth and Fatih Kaleoglu. "A Note on Copy-Protection from Random Oracles". In: *arXiv preprint arXiv:2208.12884* (2022) (cit. on p. 16).

[AKL+22] Prabhanjan Ananth, Fatih Kaleoglu, Xingjian Li, Qipeng Liu, and Mark Zhandry. *On the Feasibility of Unclonable Encryption, and More*. Cryptology ePrint Archive, Paper 2022/884. https://eprint.iacr.org/2022/884. 2022. URL: https://eprint.iacr.org/2022/884 (cit. on pp. 16, 17, 33–35).

[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 pp. 7, 16).

[AKN+23] Shweta Agrawal, Fuyuki Kitagawa, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa. "Public Key Encryption with Secure Key Leasing". In: *arXiv preprint arXiv:2302.11663* (2023) (cit. on pp. 8, 30).

[AL21] Prabhanjan Ananth and Rolando L La Placa. "Secure Software Leasing". In: *Eurocrypt* (2021) (cit. on pp. 3, 4, 6, 16).

[ALL+21] Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. "New approaches for quantum copy-protection". In: *Annual International Cryptology Conference*. Springer. 2021, pp. 526–555 (cit. on pp. 16, 33–35).

[AP08] Joel Alwen and Chris Peikert. *Generating Shorter Bases for Hard Random Lattices*. Cryptology ePrint Archive, Paper 2008/521. https://eprint.iacr.org/2008/521. 2008. URL: https://eprint.iacr.org/2008/521 (cit. on p. 57).

[Ban93] W. Banaszczyk. "New bounds in some transference theorems in the geometry of numbers." In: *Mathematische Annalen* 296.4 (1993), pp. 625–636. URL: http://eudml.org/doc/165105 (cit. on p. 20).

[BB84] C. H. Bennett and G. Brassard. "Quantum cryptography: Public key distribution and coin tossing". In: *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*. Bangalore, 1984, p. 175 (cit. on pp. 16, 17).

[BBK22] Nir Bitansky, Zvika Brakerski, and Yael Tauman Kalai. *Constructive Post-Quantum Reductions*. 2022. DOI: 10.48550/ARXIV.2203.02314. URL: https://arxiv.org/abs/2203.02314 (cit. on pp. 11, 28).

[BCM+21] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. *A Cryptographic Test of Quantumness and Certifiable Randomness from a Single Quantum Device*. 2021. arXiv: 1804.00640 [quant-ph] (cit. on p. 23).

[BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. "Factoring and pairings are not necessary for io: Circular-secure lwe suffices". In: *Cryptology ePrint Archive* (2020) (cit. on p. 16).

[BGG+14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. *Fully Key-Homomorphic Encryption, Arithmetic Circuit ABE, and Compact Garbled Circuits*. Cryptology ePrint Archive, Paper 2014/356. https://eprint.iacr.org/2014/356. 2014. URL: https://eprint.iacr.org/2014/356 (cit. on p. 8).

[BI20a] Anne Broadbent and Rabib Islam. "Quantum Encryption with Certified Deletion". In: *Lecture Notes in Computer Science* (2020), pp. 92–122. ISSN: 1611-3349. DOI: 10.1007/978-3-030-64381-2_4. URL: http://dx.doi.org/10.1007/978-3-030-64381-2_4 (cit. on p. 17).

[BI20b] Anne Broadbent and Rabib Islam. "Quantum encryption with certified deletion". In: *Theory of Cryptography Conference*. Springer. 2020, pp. 92–122 (cit. on pp. 3–5, 8, 17).

[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. 16).

[BK22] James Bartusek and Dakshita Khurana. *Cryptography with Certified Deletion*. 2022. DOI: 10.48550/ARXIV.2207.01754. URL: https://arxiv.org/abs/2207.01754 (cit. on pp. 3, 5, 17).

[BKP23] James Bartusek, Dakshita Khurana, and Alexander Poremba. *Publicly-Verifiable Deletion via Target-Collapsing Functions*. 2023. arXiv: 2303.08676 [quant-ph] (cit. on pp. 13, 55, 59).

[BL20] Anne Broadbent and Sébastien Lord. "Uncloneable Quantum Encryption via Oracles". In: *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*. Ed. by Steven T. Flammia. Vol. 158. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, 4:1–4:22. DOI: 10.4230/LIPIcs.TQC.2020.4 (cit. on pp. 3, 8, 17).

[BMSZ16] Saikrishna Badrinarayanan, Eric Miles, Amit Sahai, and Mark Zhandry. "Post-zeroizing Obfuscation: New Mathematical Tools, and the Case of Evasive Circuits". In: *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*. 2016, pp. 764–791 (cit. on p. 16).

[BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. "Pseudorandom functions and lattices". In: *Advances in Cryptology–EUROCRYPT 2012: 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings 31*. Springer. 2012, pp. 719–737 (cit. on p. 8).

[Bra18] Zvika Brakerski. *Quantum FHE (Almost) As Secure As Classical*. Cryptology ePrint Archive, Report 2018/338. https://ia.cr/2018/338. 2018 (cit. on p. 23).

[BV14a] Zvika Brakerski and Vinod Vaikuntanathan. "Efficient fully homomorphic encryption from (standard) LWE". In: *SIAM Journal on Computing* 43.2 (2014), pp. 831–871. DOI: 10.1137/120868669 (cit. on p. 8).

[BV14b] Zvika Brakerski and Vinod Vaikuntanathan. "Efficient fully homomorphic encryption from (standard) LWE". In: *SIAM Journal on computing* 43.2 (2014), pp. 831–871 (cit. on p. 6).

[CFGN96] Ran Canetti, Uri Feige, Oded Goldreich, and Moni Naor. "Adaptively Secure Multi-Party Computation". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, pp. 639–648. ISBN: 0897917855. DOI: 10.1145/237814.238015. URL: https://doi.org/10.1145/237814.238015 (cit. on p. 17).

[CJJ22] Arka Rai Choudhuri, Abhihsek Jain, and Zhengzhong Jin. "Snargs for P from LWE". In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 68–79 (cit. on p. 8).

[CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. "Hidden cosets and applications to unclonable cryptography". In: *Annual International Cryptology Conference*. Springer. 2021, pp. 556–584 (cit. on pp. 3–6, 8, 11, 16, 33).

[CMP20] Andrea Coladangelo, Christian Majenz, and Alexander Poremba. *Quantum copy-protection of compute-and-compare programs in the quantum random oracle model*. 2020. arXiv: 2009.13865 [quant-ph] (cit. on p. 16).

[CVW18] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. "GGH15 beyond permutation branching programs: proofs, attacks, and candidates". In: *Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part II 38*. Springer. 2018, pp. 577–607 (cit. on p. 16).

[DGT+10] Yevgeniy Dodis, Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. "Public-Key Encryption Schemes with Auxiliary Inputs". In: *Theory of Cryptography*. Ed. by Daniele Micciancio. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 361–381. ISBN: 978-3-642-11799-2 (cit. on pp. 11, 21, 28, 29).

[Die82] DGBJ Dieks. "Communication by EPR devices". In: *Physics Letters A* 92.6 (1982), pp. 271–272 (cit. on p. 3).

[DQV+21] Lalita Devadas, Willy Quach, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. "Succinct LWE sampling, random polynomials, and obfuscation". In: *Theory of Cryptography: 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8–11, 2021, Proceedings, Part II 19*. Springer. 2021, pp. 256–287 (cit. on p. 16).

[Gao15] Jingliang Gao. "Quantum union bounds for sequential projective measurements". In: *Physical Review A* 92.5 (2015), p. 052331 (cit. on p. 19).

[Gen09] Craig Gentry. "Fully homomorphic encryption using ideal lattices". In: *Proceedings of the forty-first annual ACM symposium on Theory of computing*. 2009, pp. 169–178 (cit. on p. 6).

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. "How to construct random functions". In: *Journal of the ACM* 33.4 (1986), pp. 792–807. ISSN: 0004-5411. DOI: 10.1145/6490.6503 (cit. on p. 73).

[GL89] O. Goldreich and L. A. Levin. "A Hard-Core Predicate for All One-Way Functions". In: *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*. STOC '89. Seattle, Washington, USA: Association for Computing Machinery, 1989, pp. 25–32. ISBN: 0897913078. DOI: 10.1145/73007.73010. URL: https://doi.org/10.1145/73007.73010 (cit. on p. 11).

[Gol06] Oded Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Cambridge University Press, 2006 (cit. on p. 76).

[Got02] Daniel Gottesman. "Uncloneable encryption". In: *arXiv preprint quant-ph/0210062* (2002) (cit. on p. 3).

[GP21] Romain Gay and Rafael Pass. "Indistinguishability obfuscation from circular security". In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 736–749 (cit. on p. 16).

[GPV07] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. *Trapdoors for Hard Lattices and New Cryptographic Constructions*. Cryptology ePrint Archive, Report 2007/432. https://eprint.iacr.org/2007/432. 2007 (cit. on pp. 4, 9, 19, 20, 22, 25, 31).

[GR02] Lov K. Grover and Terry Rudolph. "Creating superpositions that correspond to efficiently integrable probability distributions". In: *arXiv: Quantum Physics* (2002) (cit. on p. 23).

[GSW13] Craig Gentry, Amit Sahai, and Brent Waters. *Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based*. Cryptology ePrint Archive, Report 2013/340. https://ia.cr/2013/340. 2013 (cit. on pp. 6, 13, 70).

[GZ20] Marios Georgiou and Mark Zhandry. "Unclonable decryption keys". In: *Cryptology ePrint Archive* (2020) (cit. on pp. 3–5).

[HH00] L. Hales and S. Hallgren. "An improved quantum Fourier transform algorithm and applications". In: *Proceedings 41st Annual Symposium on Foundations of Computer Science*. 2000, pp. 515–525. DOI: 10.1109/SFCS.2000.892139 (cit. on p. 18).

[HMNY21a] Taiga Hiroka, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. *Certified Everlasting Zero-Knowledge Proof for QMA*. 2021. arXiv: 2109.14163 [quant-ph] (cit. on pp. 3, 5).

[HMNY21b] Taiga Hiroka, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. "Quantum Encryption with Certified Deletion, Revisited: Public Key, Attribute-Based, and Classical Communication". In: *Lecture Notes in Computer Science*. Springer International Publishing, 2021, pp. 606–636. DOI: 10.1007/978-3-030-92062-3_21. URL: https://doi.org/10.1007%2F978-3-030-92062-3_21 (cit. on p. 5).

[HMNY21c] Taiga Hiroka, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. *Quantum Encryption with Certified Deletion, Revisited: Public Key, Attribute-Based, and Classical Communication*. 2021. arXiv: 2105.05393 [quant-ph] (cit. on p. 17).

[Int15] Intercept. "How Spies Stole The Keys To The Encryption Castle". In: *https://theintercept.com/2015/02/19/great-sim-heist/*. 2015 (cit. on p. 3).

[JL00] Stanisław Jarecki and Anna Lysyanskaya. "Adaptively Secure Threshold Cryptography: Introducing Concurrency, Removing Erasures". In: *Proceedings of the 19th International Conference on Theory and Application of Cryptographic Techniques*. EUROCRYPT'00. Bruges, Belgium: Springer-Verlag, 2000, pp. 221–242. ISBN: 3540675175 (cit. on p. 17).

[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. "Indistinguishability obfuscation from well-founded assumptions". In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 60–73 (cit. on p. 16).

[KN22] Fuyuki Kitagawa and Ryo Nishimaki. "Functional Encryption with Secure Key Leasing". In: *ASIACRYPT*. 2022 (cit. on p. 5).

[KNY21a] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. "Secure Software Leasing from Standard Assumptions". In: *Theory of Cryptography*. Springer International Publishing, 2021, pp. 31–61. DOI: 10.1007/978-3-030-90459-3_2. URL: https://doi.org/10.1007%2F978-3-030-90459-3_2 (cit. on p. 6).

[KNY21b] 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 pp. 8, 16).

[LLQZ22] Jiahui Liu, Qipeng Liu, Luowen Qian, and Mark Zhandry. *Collusion Resistant Copy-Protection for Watermarkable Functionalities*. Cryptology ePrint Archive, Paper 2022/1429. https://eprint.iacr.org/2022/1429. 2022. URL: https://eprint.iacr.org/2022/1429 (cit. on pp. 3, 4).

[LZ19] Qipeng Liu and Mark Zhandry. *Revisiting Post-Quantum Fiat-Shamir*. Cryptology ePrint Archive, Paper 2019/262. https://eprint.iacr.org/2019/262. 2019. URL: https://eprint.iacr.org/2019/262 (cit. on p. 24).

[Mah18] Urmila Mahadev. *Classical Verification of Quantum Computations*. 2018. arXiv: 1804.01082 [quant-ph] (cit. on pp. 13, 70).

[MP11] Daniele Micciancio and Chris Peikert. *Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller*. Cryptology ePrint Archive, Report 2011/501. https://eprint.iacr.org/2011/501. 2011 (cit. on p. 23).

[MR04] D. Micciancio and O. Regev. "Worst-case to average-case reductions based on Gaussian measures". In: *45th Annual IEEE Symposium on Foundations of Computer Science*. 2004, pp. 372–381. DOI: 10.1109/FOCS.2004.72 (cit. on p. 20).

[MR07] Daniele Micciancio and Oded Regev. "Worst-Case to Average-Case Reductions Based on Gaussian Measures". In: *SIAM J. Comput.* 37.1 (2007), pp. 267–302. DOI: 10.1137/S0097539705447360. URL: https://doi.org/10.1137/S0097539705447360 (cit. on p. 22).

[MW05] Chris Marriott and John Watrous. *Quantum Arthur-Merlin Games*. 2005. arXiv: cs/0506068 [cs.CC] (cit. on p. 33).

[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. 17).

[Por22] Alexander Poremba. *Quantum Proofs of Deletion for Learning with Errors*. 2022. DOI: 10.48550/ARXIV.2203.01610. URL: https://arxiv.org/abs/2203.01610 (cit. on pp. 3, 5, 11, 17, 24–26).

[PR06] Chris Peikert and Alon Rosen. "Efficient Collision-Resistant Hashing from Worst-Case Assumptions on Cyclic Lattices". In: *Theory of Cryptography*. Ed. by Shai Halevi and Tal Rabin. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 145–166. ISBN: 978-3-540-32732-5 (cit. on p. 20).

[PS18] Chris Peikert and Sina Shiehian. "Privately constraining and programming PRFs, the LWE way". In: *Public-Key Cryptography–PKC 2018: 21st IACR International Conference on Practice and Theory of Public-Key Cryptography, Rio de Janeiro, Brazil, March 25-29, 2018, Proceedings, Part II 21*. Springer. 2018, pp. 675–701 (cit. on p. 77).

[Reg05] Oded Regev. "On lattices, learning with errors, random linear codes, and cryptography". In: *Journal of the ACM* 56.6 (2005), 34:1–34:40. ISSN: 0004-5411. DOI: 10.1145/1568318.1568324 (cit. on pp. 4, 8, 22, 23).

[Riv98] Ronald L. Rivest. "Can We Eliminate Certificate Revocation Lists?" In: *Proceedings Financial Cryptography '98*. FC'98 (Anguilla, British West Indies, Feb. 23–25, 1998). Ed. by Rafael Hirschfeld. Vol. 1465. Lecture Notes in Computer Science. Springer, Feb. 1998, pp. 178–183. ISBN: 978-3-540-64951-9. DOI: 10.1007/BFb0055482 (cit. on p. 3).

[SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. *Efficient Public Key Encryption Based on Ideal Lattices*. Cryptology ePrint Archive, Paper 2009/285. https://eprint.iacr.org/2009/285. 2009. URL: https://eprint.iacr.org/2009/285 (cit. on pp. 5, 23, 25).

[Stu95] S. Stubblebine. "Recent-secure authentication: enforcing revocation in distributed systems". In: *2012 IEEE Symposium on Security and Privacy*. Los Alamitos, CA, USA: IEEE Computer Society, May 1995, p. 0224 (cit. on p. 3).

[TL17] Marco Tomamichel and Anthony Leverrier. "A largely self-contained and complete security proof for quantum key distribution". In: *Quantum* 1 (July 2017), p. 14. ISSN: 2521-327X. DOI: 10.22331/q-2017-07-14-14. URL: https://doi.org/10.22331/q-2017-07-14-14 (cit. on p. 17).

[Unr13] Dominique Unruh. *Revocable quantum timed-release encryption*. Cryptology ePrint Archive, Report 2013/606. https://ia.cr/2013/606. 2013 (cit. on pp. 3, 4, 16, 17).

[Unr15] Dominique Unruh. *Computationally binding quantum commitments*. Cryptology ePrint Archive, Paper 2015/361. https://eprint.iacr.org/2015/361. 2015. URL: https://eprint.iacr.org/2015/361 (cit. on p. 24).

[Wat05] John Watrous. *Zero-knowledge against quantum attacks*. 2005. DOI: 10.48550/ARXIV.QUANT-PH/0511020. URL: https://arxiv.org/abs/quant-ph/0511020 (cit. on p. 18).

[Wie83] Stephen Wiesner. "Conjugate Coding". In: *SIGACT News* 15.1 (Jan. 1983), pp. 78–88. ISSN: 0163-5700. DOI: 10.1145/1008908.1008920. URL: https://doi.org/10.1145/1008908.1008920 (cit. on pp. 3, 16).

[Wil13] Mark M. Wilde. *Quantum Information Theory*. 1st. USA: Cambridge University Press, 2013. ISBN: 1107034256 (cit. on p. 17).

[WW21] Hoeteck Wee and Daniel Wichs. "Candidate obfuscation via oblivious LWE sampling". In: *Advances in Cryptology–EUROCRYPT 2021: 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17–21, 2021, Proceedings, Part III*. Springer. 2021, pp. 127–156 (cit. on p. 16).

[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. 3).

[Zha20] Mark Zhandry. *Schrödinger's Pirate: How To Trace a Quantum Decoder*. Cryptology ePrint Archive, Paper 2020/1191. https://eprint.iacr.org/2020/1191. 2020. URL: https://eprint.iacr.org/2020/1191 (cit. on pp. 33–35).

[Zha21] Mark Zhandry. "Quantum Lightning Never Strikes the Same State Twice. Or: Quantum Money from Cryptographic Assumptions". In: *J. Cryptol.* 34.1 (Jan. 2021). ISSN: 0933-2790. DOI: 10.1007/s00145-020-09372-x. URL: https://doi.org/10.1007/s00145-020-09372-x (cit. on p. 8).