



Quantum Key-Revocable Dual-Regev Encryption, Revisited

Prabhanjan Ananth¹ , Zihan Hu² , and Zikuan Huang³ 

¹ UCSB, Santa Barbara, USA

prabhanjan@cs.ucsb.edu

² EPFL, Lausanne, Switzerland

zihan.hu@epfl.ch

³ Tsinghua University, Beijing, China

hzk21@mails.tsinghua.edu.cn

Abstract. Quantum information can be used to achieve novel cryptographic primitives that are impossible to achieve classically. A recent work by Ananth, Poremba, Vaikuntanathan (TCC 2023) focuses on equipping the dual-Regev encryption scheme, introduced by Gentry, Peikert, Vaikuntanathan (STOC 2008), with key revocation capabilities using quantum information. They further showed that the key-revocable dual-Regev scheme implies the existence of fully homomorphic encryption and pseudorandom functions, with both of them also equipped with key revocation capabilities. Unfortunately, they were only able to prove the security of their schemes based on new conjectures and left open the problem of basing the security of key revocable dual-Regev encryption on well-studied assumptions.

In this work, we resolve this open problem. Assuming polynomial hardness of learning with errors (over sub-exponential modulus), we show that key-revocable dual-Regev encryption is secure. As a consequence, for the first time, we achieve the following results:

- Key-revocable public-key encryption and key-revocable fully-homomorphic encryption satisfying classical revocation security and based on polynomial hardness of learning with errors. Prior works either did not achieve classical revocation or were based on sub-exponential hardness of learning with errors.
- Key-revocable pseudorandom functions satisfying classical revocation from the polynomial hardness of learning with errors. Prior works relied upon unproven conjectures.

1 Introduction

Leveraging fundamental principles of quantum information to achieve cryptographic notions, that are otherwise impossible to achieve classically, is an exciting research direction. In the past few years, a dizzying variety of quantum cryptographic primitives, termed as *unclonable* primitives, have been studied. Underlying the unclonable primitives is the no-cloning principle of quantum mechanics [WZ82, Die82] which states that quantum states, unlike classical strings, cannot be copied. The recent surge in the development of unclonable primitives has

resulted in innovative approaches to tackle many real-world security challenges, including protection against anti-piracy [Aar09], privacy concerns in blockchain technology [AGKZ20], and provable deletion of cryptographic data from the web [BI20, BL20].

We focus on the task of securely leasing or revoking cryptographic keys using the tools of quantum information. Before precisely stating the problem that we set out to address, let us consider two scenarios: (a) Imagine a manager needing to temporarily delegate their duties, including access to sensitive encrypted data, to their subordinate by sharing cryptographic keys. The challenge is ensuring the subordinate's access is revoked upon the manager's return, a task that is impossible to achieve with classical keys, (b) If a cryptographic key is stolen from a device, unless the attacker has left a trace, it becomes challenging to detect such an attack and report it.

Quantum information presents a unique approach of tackling both of the above aforementioned problems.

OUR FOCUS. A major focus of our work is on protecting decryption keys. Specifically, we focus on the popular dual-Regev public-key encryption scheme of [GPV08] (also, referred to as the GPV encryption scheme), which has inspired the design of many lattice-based cryptographic primitives [BGG+14, Mah18, BDGM20, Qua20]. A key-revocable dual-Regev public-key encryption scheme, first introduced in [APV23], is the same as the dual-Regev scheme except that we have the additional guarantee that the decryption keys can alternately be represented as quantum states. Any user in possession of the quantum decryption key can decrypt ciphertexts just the way he would have been able to do if he had a classical decryption key. The security guarantee stipulates that once the user returns the quantum decryption key, they will lose the ability to decrypt ciphertexts and in particular, we require that the semantic security of dual-Regev encryption still hold. We refer the reader to Sect. 1.1 for a more detailed description of the key-revocable dual Regev public-key encryption scheme.

KEY-REVOCABLE SECURITY OF DUAL-REGEV: MOTIVATION. Proving the security of key-revocable dual-Regev encryption could lead to adding key revocation capabilities to other cryptographic primitives. Indeed, [APV23] showed that key-revocable dual-Regev encryption can be leveraged to prove the existence of fully homomorphic encryption and pseudorandom functions equipped with key revocation capabilities. The structure of dual-Regev encryption was crucially exploited in these applications.

There is also an aesthetic reason behind studying this problem. Dual-Regev public-key encryption is an elegant construction that is taught in most graduate classes on lattice-based cryptography. Understanding whether it satisfies key-revocable security is a natural theoretical question.

The work of [APV23] attempted to prove the key-revocable security of dual-Regev encryption. Unfortunately, they were only able to prove the security of this construction based on a new unfounded conjecture. They leave the problem of proving the key-revocable security of dual Regev encryption on concrete computational assumptions as an important open problem. In this same

work, inspired by the literature on certified deletion [BI20, HMNY21, BK22], they define a stronger property called classical revocation: instead of the user being asked to return the state, they are only asked to return a classical string that certifies that the quantum decryption key has been deleted. After the state has been deleted, as before, we require the semantic security of dual-Regev encryption to still hold. [APV23] relied upon yet another new conjecture to show that dual-Regev encryption satisfied classical key-revocation security. The reliance on both these conjectures makes the current state of affairs rather unsatisfactory. [APV23] left open the problem of basing key-revocation security of dual-Regev encryption on well-studied cryptographic assumptions.

Main Result. In this work, we resolve this open problem. We show the following:

Theorem 1. *Assuming polynomial hardness of learning with errors over sub-exponential modulus¹, dual-Regev encryption is key-revocable. Moreover, this scheme satisfies the classical revocation property.*

Applications. By combining the above theorem with the applications of key-revocable dual-Regev encryption in [APV23], we obtain the following results:

MAIN APPLICATION: We present the first result of *key-revocable pseudorandom functions* based on the polynomial hardness of learning with errors and also simultaneously satisfies classical revocation property. Prior work by [APV23] relied upon unproven conjectures.

OTHER APPLICATIONS: We also achieve other applications that are in some aspects better than the previous works.

1. We present the first result of *key-revocable public-key encryption* that is based on polynomial hardness of learning with errors and simultaneously satisfies classical revocation property. Prior works by [AKN+23, CGJL23] satisfied one but not the other.
2. We present the first result of *key-revocable fully homomorphic encryption* that is based on polynomial hardness of learning with errors and simultaneously satisfies classical revocation property. Prior work by [CGJL23] achieved this result from sub-exponential hardness of learning with errors.

MAIN TECHNICAL CONTRIBUTION: At the heart of our result is a new search-to-decision reduction that reduces a quantum distinguisher that breaks the semantic security of dual-Regev encryption into a quantum adversary that can solve an inhomogeneous short integer solution (ISIS) problem. Our search-to-decision reduction is qualitatively different from [APV23] who rely upon Goldreich-Levin reduction over large finite fields. In addition to the fact that [APV23] relies

¹ By aggressively setting the parameters, it would suffice to just assume polynomial hardness of learning with errors over quasi-polynomial modulus.

upon a conjecture, their reduction necessarily² incurs a loss that is inversely proportional to q , where q is the size of the field. Since they need to set q to be sub-exponential in the security parameter, this means that their reduction suffers from sub-exponential loss. On the other hand, our ISIS solver only incurs inverse polynomial loss, independent of q .

RELATED WORKS: It would be remiss not to discuss two other related prior works.

Chardouvelis, Goyal, Jain, Liu [CGJL23] present instantiations of key-revocable public-key encryption and fully homomorphic encryption. Moreover, their schemes satisfy classical key-revocation security³. *There are two advantages of our work over theirs:*

- *They do not have any results on pseudorandom functions,*
- *They assume sub-exponential hardness of learning with errors whereas we only assume polynomial hardness of learning with errors.*

Besides that, our work **fundamentally** differs from their work, both in terms of constructions and its analysis. Let us begin by highlighting the differences in the construction.

At a high level, our construction is the same as the dual-Regev public-key encryption scheme except for the quantum decryption key, whereas [CGJL23] builds a new encryption scheme inspired by noisy trapdoor claw-free functions (NTCF) introduced by [BCM+21]. Specifically, they repeat many instantiations of NTCFs in parallel and use that to build a key-revocable public-key encryption scheme. The NTCF itself is instantiated using the (original) Regev public-key encryption scheme.

Even the overall approach in the analysis is quite different: we do a reduction from decision LWE to SIS whereas they do a search-to-decision reduction for LWE itself. In the analysis, we use Gaussian collapsing lemma [Por22] and introduce a new lemma, lemma 6 which is distinct in our work. Given the fact that the constructions are very different, unsurprisingly, the implementation details also vary quite a bit in both the works. For instance, since they do parallel repetition, their extraction method is more complicated since the adversary could have broken any one of the instantiations.

Agrawal, Kitagawa, Nishimaki, Yamada, Yamakawa [AKN+23] present an instantiation of key-revocable public-key encryption based on the existence of any post-quantum secure public-key encryption scheme. They also present other key-revocable notions, such as functional encryption, that are not covered in this work. *There are two advantages of our work over theirs:*

- *They do not prove the classical key revocation security of their scheme,*

² Their starting point is the classical Goldreich Levin reduction over finite fields by Dodis et al. [DGT+10]. This reduction already suffers from a loss that is inversely proportional to q .

³ In fact, they satisfy a much stronger property where the communication with the user can be completely classical.

- They also do not provide any positive results on either fully homomorphic encryption or pseudorandom functions.

Both the works, [CGJL23] and [AKN+23], come up with arguably more involved constructions of key-revocable public-key encryption which make it unwieldy to extend their techniques to get new applications.

1.1 Technical Overview

In this section, we give an overview of the main ideas and techniques underlying our proofs.

Key-Revocable Dual-Regev Public-Key Encryption. We first recall the key-revocable dual-Regev constructions from [APV23]. This part has been reproduced verbatim from their work.

- **KeyGen**(1^λ): Sample a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a short trapdoor basis $\text{td}_{\mathbf{A}}$ for it. The (quantum) decryption key is a Gaussian superposition of ISIS solutions, which is generated by the following procedure: Create a Gaussian superposition of short vectors \mathbf{x} , compute the image $\mathbf{A} \cdot \mathbf{x} \pmod{q}$ in the second register to get

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

where $\rho_\sigma(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/\sigma^2)$ is the Gaussian measure, for some $\sigma > 0$, and measure the second register to the Gaussian coset state

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

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

Finally we set $\text{PK} = (\mathbf{A}, \mathbf{y})$, $\text{MSK} = \text{td}_{\mathbf{A}}$ and $\rho_{\text{SK}} = |\psi_{\mathbf{y}}\rangle$.

- **Enc**(PK, μ): To encrypt a bit $\mu \in \{0, 1\}$, sample a random string $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^m$ together with discrete Gaussian errors $\mathbf{e} \in \mathbb{Z}^m$ and $e' \in \mathbb{Z}$, and output a classical ciphertext CT given by

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

- **Dec**($\rho_{\text{SK}}, \text{CT}$): First apply the unitary $U : |\mathbf{x}\rangle|0\rangle \rightarrow |\mathbf{x}\rangle|\text{CT} \cdot (-\mathbf{x}, 1)^\top\rangle$ on input $\rho_{\text{SK}} \otimes |0\rangle\langle 0|$, and then measure the second register in the computational basis. Because ρ_{SK} is supposed to be the Gaussian coset state $|\psi_{\mathbf{y}}\rangle$, which is a superposition of short vector \mathbf{x} subject $\mathbf{A} \cdot \mathbf{x} = \mathbf{y}$, we obtain an approximation of $\mu \cdot \left\lfloor \frac{q}{2} \right\rfloor$ from which we can recover μ .

- $\text{Revoke}(\text{PK}, \text{MSK}, \rho)$: Apply the projective measurement $\{|\psi_y\rangle\langle\psi_y|, I - |\psi_y\rangle\langle\psi_y|\}$ onto ρ using the master secret key td_A ⁴. Output **Valid** if the measurement succeeds, and output **Invalid**, otherwise.

Consider an efficient adversary \mathcal{A} . It receives as input a state $|\psi_y\rangle$ from the challenger and computes a state $\rho_{R, \text{Aux}}$ on two registers R and Aux . Subsequently, the adversary returns system R to the challenger, while retaining system Aux as quantum advice for subsequent steps. Informally speaking, we say that the above scheme is secure if \mathcal{A} wins both of the following events simultaneously only with negligible probability:

- Revoke on the system R outputs **Valid**.
- Using Aux , \mathcal{A} can distinguish $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{s}^\top \mathbf{y} + e' + \lfloor \frac{q}{2} \rfloor)$ versus $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{s}^\top \mathbf{y} + e')$

Starting Point. Inspired by [APV23], we undertake the following approach. Suppose there did exist an efficient adversary \mathcal{A} that is successful in violating the security of the above construction. We reduce \mathcal{A} into an SIS solver \mathcal{B} , which is described as follows: it first runs \mathcal{A} on input $(\mathbf{A}, \mathbf{y}, |\psi_y\rangle)$ to obtain a state ρ on two registers R and Aux . Then, \mathcal{B} needs to be cleverly designed in such a way that it recovers a short vector \mathbf{x}_0 from R and a short vector \mathbf{x}_1 from Aux satisfying the following properties:

- $\mathbf{A}\mathbf{x}_0 = \mathbf{y}$, $\mathbf{A}\mathbf{x}_1 = \mathbf{y}$ and,
- $\mathbf{x}_0 \neq \mathbf{x}_1$.

Once both the vectors \mathbf{x}_0 and \mathbf{x}_1 are recovered then it simply sets the SIS solution to be $\mathbf{x}_0 - \mathbf{x}_1$.

While [APV23] set out on this route, they only managed to show such a reduction based on a new conjecture. The core reason behind this is the fact that it is challenging to be able to *simultaneously* recover two *distinct* short solutions from *two potentially entangled* registers R and Aux . An attempt to recover \mathbf{x}_0 from R could invariably disturb the part of the state on Aux such that it is no longer possible to recover \mathbf{x}_1 . Any approach we undertake should tackle this challenge.

Our Approach. We propose a three-step approach to prove the security of key-revocable dual-Regev encryption based on learning with errors.

- **STEP 1.** In the first step, we transform the intermediate state ρ (on R and Aux) produced by \mathcal{A} into a “good state” ρ_{good} . This step doesn’t need to always succeed. We require two guarantees here: (a) this step aborts with probability bounded away from 1 and, (b) conditioned on not abort, the output of this step is a good state ρ_{good} such that the revocation on R succeeds with non-negligible probability and Step 2 works.

⁴ [APV23] showed how to implement this projective measurement efficiently with the trapdoor td_A .

- STEP 2. Suppose the output of Step 1 is ρ_{good} . We require that as long as ρ_{good} is a good state then, from AUX, we should be able to recover a short vector \mathbf{x}_1 such that $\mathbf{A}\mathbf{x}_1 = \mathbf{y}$. More importantly, we should be able to recover \mathbf{x}_1 with overwhelming probability.
- STEP 3. We recover a short vector \mathbf{x}_0 from the register R such that $\mathbf{A}\mathbf{x}_0 = \mathbf{y}$. Our hope is that \mathbf{x}_0 and \mathbf{x}_1 are distinct and if this is the case then $\mathbf{x}_0 - \mathbf{x}_1$ is a non-trivial short solution in the kernel of \mathbf{A} .

The easiest step to realize is Step 3. Suppose we have the guarantee that we can recover \mathbf{x}_1 from AUX with overwhelming probability. By invoking *almost as good as new* lemma [Aar16], we can show that the state ρ after Step 2 is not disturbed by much. This means that **Revoke** still succeeds on R with inverse polynomial probability. This further implies that measuring the register R yields a short vector \mathbf{x}_0 . Then using a simplified analysis of [APV23], we can argue that $\mathbf{x}_0 \neq \mathbf{x}_1$, completing the proof.

We focus our attention on implementing Steps 1 and 2. Our main technical contribution will lie in Step 2.

IMPLEMENTING STEP 1: To implement Step 1, we rely upon the threshold implementation technique introduced by Zhandry [Zha20]. Threshold implementation is a technique employed to get an estimate of the success probability of a POVM on a state. In our context, we employ this technique to test whether the adversary acting upon AUX register of ρ is successful in violating the security of key-revocable dual-Regev encryption scheme. Formally, we define the threshold implementation operator $\text{TI}_{\frac{1}{2}+\gamma}$, where γ is some inverse polynomial, with the following properties:

1. $\text{TI}_{\frac{1}{2}+\gamma}$ is akin to a projector-like operator, collapsing the state to a γ -good state ρ_{good} capable of distinguishing between $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{s}^\top \mathbf{y} + e')$ and (\mathbf{u}, r) for \mathbf{u}, r being sampled uniformly randomly with probability 2γ (referred to as a “ γ -good state”) when $\text{TI}_{\frac{1}{2}+\gamma}$ outputs 1, or to some other state when $\text{TI}_{\frac{1}{2}+\gamma}$ outputs 0.
2. For a successful adversary, applying $\text{TI}_{\frac{1}{2}+\gamma}$ with an inverse polynomial γ on ρ_{AUX} results in an output of 1 with noticeable probability.
3. Upon applying $\text{TI}_{\frac{1}{2}+\gamma}$ again on a γ -good state, it yields an output of 1 with probability 1.

To summarize, as long as \mathcal{A} is a successful adversary, $\text{TI}_{\frac{1}{2}+\gamma}$ collapses ρ into a good state ρ_{good} with inverse polynomial probability.

IMPLEMENTING STEP 2: As mentioned earlier, implementing Step 2 is our main technical contribution.

It was already shown by [APV23] that \mathbf{x}_1 can be extracted from AUX. However, the success probability of their extraction mechanism was only inverse polynomial which is insufficient for our purpose. Instead, we completely depart

from [APV23] and propose a novel extraction method. This high-level approach is inspired by [CGJL23] although they study for a completely different construction.

At a high level, our extractor proceeds by guessing each entry of \mathbf{x}_1 , where \mathbf{x}_1 is a short solution mapping \mathbf{A} to \mathbf{y} , one coordinate at a time. For each coordinate, we try all possible values and using the distinguisher, test which of our guesses was correct. Recall that there are exponentially many short vectors that map \mathbf{A} to \mathbf{y} . But once we apply the Gaussian collapsing lemma [Por22, LMZ23], we can replace the state $|\psi_{\mathbf{y}}\rangle$ with $|\mathbf{x}_1\rangle$. While recovering, say, the i^{th} coordinate of \mathbf{x}_1 , we use the distinguisher on AUX to figure out whether the guess for the i^{th} coordinate was correct or not. However, this has to be handled with care. Since the distinguisher has quantum auxiliary advice, we cannot keep hoping to run the distinguisher again and again. After the first run, the state of the distinguisher could be damaged making it useless for future iterations. So we need to come up with a mechanism to check if a guess is correct or not while maintaining the quantum state. Making crucial use of threshold implementation along with techniques from lattice-based cryptography, we show how to implement this.

Our extractor is described as follows:

1. Initialize $\mathbf{x} = \mathbf{0}$ as the output register.
2. For each position $i \in [m]$ and each guess g_i , we test whether the i -th entry \mathbf{x}_1 is g_i by:
 - (a) Applying $\text{TI}_{\frac{1}{2}+\gamma'}(i, g_i)$ on system AUX , where $\text{TI}_{\frac{1}{2}+\gamma'}(i, g_i)$ is a threshold implementation that ‘tests’ whether the state is γ' -good at distinguishing between $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top + c \cdot \hat{\mathbf{i}}, \mathbf{s}^\top \mathbf{y} + c \cdot g_i + e')$ (where $c \xleftarrow{\$} \mathbb{Z}_q$ and $\hat{\mathbf{i}}$ is the unit vector on the i -th dimension) and (\mathbf{u}, r) (where $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m, r \xleftarrow{\$} \mathbb{Z}_q$).
 - (b) If the output is 1, set $\mathbf{x}_i = g_i$.
 - (c) If the output is 0, skip to the next iteration.
3. Output \mathbf{x} .

We argue that our extractor outputs \mathbf{x}_1 with nearly perfect probability if $\text{TI}_{\frac{1}{2}+\gamma}$ on ρ_{AUX} outputs 1. Zhandry [Zha20] demonstrates that for two threshold implementations concerning computationally indistinguishable tasks (e.g., distinguishing $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{s}^\top \mathbf{y} + e')$ from (\mathbf{u}, r) , and distinguishing $(\mathbf{u}, \mathbf{u}^\top \mathbf{x}_1 + e')$ from (\mathbf{u}, r)), their outputs are closely related. Now, considering each guess g_i for position i :

- If the guess is correct (i.e., the i -th entry of \mathbf{x}_1 is g_i), the distribution $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top + c \cdot \hat{\mathbf{i}}, \mathbf{s}^\top \mathbf{y} + c \cdot g_i + e')$ is computationally indistinguishable from the distribution $(\mathbf{u}, \mathbf{u}^\top \mathbf{x}_1 + e')$, and thus also from $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{s}^\top \mathbf{y} + e')$. Given ρ'_{AUX} is a γ -good state, $\text{TI}_{\frac{1}{2}+\gamma'}(i, g_i)$ outputs 1 with $1 - \text{negl}$ probability if all other threshold implementations are ignored (i.e., applied $\text{TI}_{\frac{1}{2}+\gamma'}(i, g_i)$ just after $\text{TI}_{\frac{1}{2}+\gamma}$).
- If the guess is incorrect, the distribution $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top + c \cdot \hat{\mathbf{i}}, \mathbf{s}^\top \mathbf{y} + c \cdot g_i + e')$ is computationally indistinguishable from (\mathbf{u}, r) . Consequently, any state provides no advantage as advice, and $\text{TI}_{\frac{1}{2}+\gamma'}$ outputs 1 with negl probability.

Finally, we apply the quantum union bound to all measurements to demonstrate that the probability of no error occurring during our testing procedure is $1 - \text{negl}$.

In the above proof, we omitted a major issue. Recall that in Step 1, we implement threshold implementation to project the state ρ onto a good state ρ_{good} . Moreover, this threshold implementation is designed to check if the adversary can distinguish between the distributions $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{s}^\top \mathbf{y} + e')$ and (\mathbf{u}, r) . As discussed above, at some point, in the intermediate hybrids we need to change these distributions. Once we switch the distributions, the threshold implementation might only work with negligible probability. Our hope, in some cases invoking learning with errors, is to argue that this does not happen. However, it is not clear how to carry out this reduction. After all, the threshold implementation as defined by [Zha20] operates on a superposition of exponentially many samples from a distribution and so, given just one sample from a distribution, it is not possible to perform threshold implementation. We present a useful lemma (in Sect. 5) where we argue that operationally, the guarantees of threshold implementation (including the output and the residual state) are not affected when one distribution is replaced with another computationally indistinguishable distribution.

2 Preliminaries

We use standard notations throughout this work. We assume that the reader is familiar with quantum computing and lattices. We refer the reader to the full version [AHH24] for a complete presentation of this section.

2.1 Quantum Computing

We will use the following lemma.

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

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

2.2 Lattices and Cryptography

We adapt notations from [APV23] and keep it same as much as we can. The following subsection is copied verbatim from [APV23].

In this work, we mainly consider *q-ary lattices* Λ 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 $\mathbf{A}\mathbf{c} = \mathbf{y} \pmod{q}$.

Definition 1 (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

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

where $\rho_\sigma(\mathbf{x}) := \exp(-\pi\|\mathbf{x}\|^2/\sigma^2)$ is the Gaussian distribution.

We will use the following two results.

Lemma 2 (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 3 (min-entropy of the truncated discrete Gaussian, [APV23], 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

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

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

2.3 Threshold Implementation and Its Approximate Version

In the subsection, we review some techniques called *Threshold Implementation* [ALL+21], which is a simple extension of Projective Implementation [Zha20].

Theorem 3 (Threshold implementation, [ALL+21]). *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 $(\mathsf{TI}_{\gamma}(\mathcal{P}), \mathbf{I} - \mathsf{TI}_{\gamma}(\mathcal{P}))$ 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 $\lambda_i \leq \gamma$.
- $\mathbf{I} - \mathsf{TI}_{\gamma}(\mathcal{P})$ projects a quantum state into the subspace spanned by $\{|\psi_i\rangle\}$ whose eigenvalues λ_i satisfy the property $\lambda_i > \gamma$.

Unfortunately, the threshold implementation can, in general, not be efficiently computable. However, inspired by the work of Marriott and Watrous [MW05], Zhandry [Zha20] showed that the approximate version of the threshold implementation can be implemented efficiently as long as the POVM is a mixture of projective measurements. We first review the definition of mixture of projective measurements.

Definition 2 (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 index set \mathcal{I} . Then, $\mathcal{P}_D = (P_D, Q_D)$ is the following mixture of projective measurements:*

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

In other words, \mathcal{P}_D is the same as first sampling i according to the distribution D , and then applying the projective measurements \mathcal{P}_i .

For any mixture of projective measurements \mathcal{P}_D , the approximate threshold implementation satisfies the following properties.

Lemma 4 (Approximate threshold implementation, Theorem 6.2 in [Zha20] and Corollary 1 in [ALL+21]). 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 $\text{ATI}_{\mathcal{P}, D, \gamma}^{\varepsilon, \delta}$, interpreted as the POVM element corresponding to outcome 1, such that the following holds:

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

3 Definition: Key-Revocable Public-Key Encryption

A key-revocable public-key encryption is a type of public-key encryption. Consider the case where the secret key holder wishes to temporarily give the secret key to a third party and later wants to take it back while maintaining the security i.e. the third party upon taken its key away, can't decrypt any message later. This is impossible in the classical case since the third party can always copy the secret key locally. But we may achieve this functionality by representing the secret key as a quantum state.

Definition 3 (Key-Revocable Public-Key Encryption [APV23]). A key-revocable public-key encryption scheme consists of efficient algorithms $(\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$, where Enc is a PPT algorithm and $\text{KeyGen}, \text{Dec}, \text{Revoke}$ are QPT algorithms defined as follows:

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

Correctness of Decryption. For $\mu \in \{0, 1\}$, the following holds:

$$\Pr \left[\mu \leftarrow \text{Dec}(\rho_{\text{SK}}, \text{CT}) : \begin{array}{l} (\text{PK}, \text{MSK}, \rho_{\text{SK}}) \leftarrow \text{KeyGen}(1^\lambda) \\ \text{CT} \leftarrow \text{Enc}(\text{PK}, \mu) \end{array} \right] \geq 1 - \text{negl}.$$

Correctness of Revocation. The following holds:

$$\Pr [\text{Valid} \leftarrow \text{Revoke}(\text{PK}, \text{MSK}, \rho_{\text{SK}}) : (\text{PK}, \text{MSK}, \rho_{\text{SK}}) \leftarrow \text{KeyGen}(1^\lambda)] \geq 1 - \text{negl}.$$

3.1 Security Definition

The security captures the case where the adversary is given the key and later taken back. After that, if the key passes the revocation check the adversary is asked to play a CPA like game that it is given either the ciphertext of a chosen message or a random message. The adversary wins if it can distinguish between these two cases.

Definition 4. A key-revocable public-key encryption scheme $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ is (ϵ, δ) -secure if, for every QPT adversary \mathcal{A} with

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

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

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

where $\text{Expt}^{\Sigma, \mathcal{A}}(1^\lambda, b)$ is defined as Fig. 1.

If $\delta(\lambda) = 1 - \frac{1}{\text{poly}(\lambda)}$ and $\epsilon(\lambda) = \text{negl}(\lambda)$, we simply say the key-revocable encryption scheme is secure.

$\text{Expt}^{\Sigma, \mathcal{A}}(1^\lambda, b)$:

Initialization Phase:

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

Revocation Phase:

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

Guessing Phase:

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

Fig. 1. Security Experiment

4 Construction: Key Revocable Dual-Regev Encryption

The construction is exactly the same as the construction in [APV23]. We include the construction here for completeness.

Construction 4 (Key Revocable Dual-Regev Encryption [APV23]). *Let $n, m \in \mathbb{N}$ and $q \geq 2$ be a prime, each parameterized by $\lambda \in \mathbb{N}$. Let $\alpha, \beta, \sigma > 0$ be parameters. The key-revocable public key scheme $\text{RevDual} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ consists of the following QPT algorithms:*

- $\text{KeyGen}(1^\lambda) \rightarrow (\text{PK}, \rho_{\text{SK}}, \text{MSK})$: Sample $(\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \text{td}_{\mathbf{A}}) \leftarrow \text{GenTrap}(1^n, 1^m, q)$ where GenTrap is the algorithm that generates the LWE matrix with its trapdoor. Then generate a Gaussian superposition $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \text{GenGauss}(\mathbf{A}, \sigma)$ ⁵ for some $\mathbf{y} \in \mathbb{Z}_q^n$. Output $\text{PK} = (\mathbf{A}, \mathbf{y})$, $\rho_{\text{SK}} = |\psi_{\mathbf{y}}\rangle$ and $\text{MSK} = \text{td}_{\mathbf{A}}$.
- $\text{Enc}(\text{PK}, \mu) \rightarrow \text{CT}$: to encrypt a bit $\mu \in \{0, 1\}$, sample a random vector $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ and errors $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$ and $\mathbf{e}' \sim D_{\mathbb{Z}, \beta q}$ and output the ciphertext pair $\text{CT} = (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top \pmod{q}, \mathbf{s}^\top \mathbf{y} + \mathbf{e}' + \mu \cdot \lfloor \frac{q}{2} \rfloor \pmod{q}) \in \mathbb{Z}_q^m \times \mathbb{Z}_q$.
- $\text{Dec}(\rho_{\text{SK}}, \text{CT}) \rightarrow \{0, 1\}$: to decrypt CT , apply the unitary $U : |\mathbf{x}\rangle|0\rangle \rightarrow |\mathbf{x}\rangle|\text{CT} \cdot (-\mathbf{x}, 1)^\top\rangle$ on input $|\psi_{\mathbf{y}}\rangle|0\rangle$, where $\rho_{\text{SK}} = |\psi_{\mathbf{y}}\rangle$, 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.
- $\text{Revoke}(\text{MSK}, \text{PK}, \rho) \rightarrow \{\top, \perp\}$: on input $\text{td}_{\mathbf{A}} \leftarrow \text{MSK}$ and $(\mathbf{A}, \mathbf{y}) \leftarrow \text{PK}$, 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 $\text{QSampGauss}(\mathbf{A}, \text{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ ⁶. Output \top if the measurement is successful, and \perp otherwise.

From [APV23], this construction satisfies the correctness of decryption and the correctness of revocation. In this work, we will focus on showing the construction is in fact secure.

Theorem 5. *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 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 polynomial hardness of $\text{LWE}_{n, q, \alpha q}^m$ with sub-exponential modulus, the scheme $\text{RevDual} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ in Construction 4 is a secure key-revocable public-key encryption scheme according to Definition 4.*

We organize the proof of Theorem 5 in the following way:

- In Sect. 5, we prove an important property for approximate threshold implementation, which allows us to do hybrid arguments between approximate threshold implementation on computationally indistinguishable distributions.
- In Sect. 6, we present our construction for the almost perfect preimage extractor that lies in the heart of our result.
- In Sect. 7, we complete our proof of the above theorem.

⁵ The detailed description of GenGauss can be found in [APV23].

⁶ The detailed description of QSampGauss can be found in [APV23].

5 Indistinguishability on Approximate Threshold Implementation

Zhandry [Zha20] analyzed the relationship between the output distribution of $\text{TI}_{\gamma_0}(\mathcal{P}_{D_0})$ and $\text{TI}_{\gamma_1}(\mathcal{P}_{D_1})$ (and $\text{ATI}_{\mathcal{P}, D_0, \gamma_0}$ and $\text{ATI}_{\mathcal{P}, D_1, \gamma_1}$) for some thresholds γ_0 and γ_1 on the same state for two computationally indistinguishable distributions D_0 and D_1 . However, in our work, we also care about the residual state after applying the procedures. So we give a more precise analysis below.

In this section, we show how to leverage a (possibly not efficiently constructible) quantum state ρ on which $\text{ATI}_{\mathcal{P}, D_0, \gamma}$ and $\text{ATI}_{\mathcal{P}, D_1, \gamma}$ behave differently to construct a QPT distinguisher (with auxiliary state ρ) for D_0 and D_1 . This can be viewed as an extension of Theorem 6.5 and Corollary 6.9 in [Zha20].

This result allows us to do hybrid arguments between $\text{ATI}_{\mathcal{P}, D_0, \gamma}$ and $\text{ATI}_{\mathcal{P}, D_1, \gamma}$ with exactly the same threshold parameter γ for computationally indistinguishable distributions D_0 and D_1 even when an efficient quantum procedure is applied on the residual state after ATI. Notably, it applies even when we need some classical advice to sample from D_0 and D_1 , in which case, our QPT distinguisher additionally takes the same classical advice and distinguishes D_0 and D_1 .

Lemma 5. *Let \mathcal{P} be a collection of projective measurements indexed by some set \mathcal{I} . Suppose \mathcal{P} can be implemented by a quantum circuit of size $|\mathcal{P}|$. Let D_0, D_1 be two efficiently sampleable distributions over \mathcal{I} . For any state $\rho \in \mathcal{D}(\mathcal{H})$, denote $(b, \rho') \leftarrow \text{ATI}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}(\rho)$ be the procedure that runs $\text{ATI}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$ on state ρ , and gets an output b and the post-measurement state ρ' . For any polynomial μ , any quantum state ρ and any (possibly quantum) predicate $h : \{0, 1\} \times \mathcal{D}(\mathcal{H}) \rightarrow \{0, 1\}$ with circuit size $|h|$, if*

$$\left| \Pr[h(b, \rho') = 1 | (b, \rho') \leftarrow \text{ATI}_{\mathcal{P}, D_0, \gamma}^{\epsilon, \delta}(\rho)] - \Pr[h(b, \rho') = 1 | (b, \rho') \leftarrow \text{ATI}_{\mathcal{P}, D_1, \gamma}^{\epsilon, \delta}(\rho)] \right| \geq \frac{1}{\mu(\lambda)}.$$

Then there exists a quantum circuit \mathcal{C} of size $\text{poly}(\lambda, 1/\epsilon, \log(1/\delta), \mu, |\mathcal{P}|, |h|)$ (which only uses the quantum circuits to implement \mathcal{P} , h and to sample D_0, D_1 as a black box) such that

$$\left| \Pr[\mathcal{C}(\rho, x) = b : \stackrel{\$}{\underset{x \sim D_b}{\leftarrow}} \{0, 1\}] - \frac{1}{2} \right| \geq \frac{1}{(\mu(\lambda))^3 \cdot \text{poly}(\lambda, 1/\epsilon, \log(1/\delta))}$$

which is an inverse polynomial if μ is a polynomial.

Proof. The proof follows the same idea as the proof for Theorem 6.5 in [Zha20]. Roughly speaking, the output of $\text{ATI}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$ can be approximated up to inverse polynomial additive error given only polynomial samples from D . We refer the reader to the full version [AHH24] for the full proof.

ATI may change the input state in an essential way even when it outputs 1 with overwhelming probability because ATI is not a projector. For example, let a pure quantum state ρ be a superposition of eigenvectors (of \mathcal{P}_D) $|\psi_i\rangle$ whose

eigenvalues λ_i satisfy the property $\lambda_i \geq \gamma + 10\epsilon$. If we apply $\text{ATI}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$ on ρ , we will get outcome 1 with almost certainty, but the residual state ρ' may lose coherence and become closer to a mixture of $|\psi_i\rangle$.

When we know the ATI outputs 0 or 1 with overwhelming probability, it is a good idea to *minimize the disturbance* by purifying ATI and performing uncomputation, just like the famous gentle measurements. To be more precise, we consider the projective version of $\text{ATI}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$. Formally, $\text{ATI}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$ can be written as introducing $\text{poly}(1/\epsilon, \log(1/\delta))$ ancillas initialized as $|0\rangle$, applying a unitary U on the state, and then applying a projective measurement $(|0\rangle\langle 0|, |1\rangle\langle 1|)$ on the output register of state to get the output. We will denote the binary-outcome projective measurement $(U^\dagger|0\rangle\langle 0|U, U^\dagger|1\rangle\langle 1|U)$ as $\overline{\text{ATI}}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$, the projective version of $\text{ATI}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$, which also has size $\text{poly}(1/\epsilon, \log(1/\delta))$. By definition, for any quantum state ρ , the output distribution of running $\overline{\text{ATI}}$ on ρ along with enough fresh ancillas is the same as the output distribution of running ATI on ρ (but the residual states are different).

Roughly speaking, $\overline{\text{ATI}}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$ does the same thing as $\text{ATI}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$ except that it uncomputes intermediate results. Notice that a quantum query to function f is implemented as $U_f : |x\rangle|y\rangle \rightarrow |x\rangle|y \oplus f(x)\rangle$, whose inverse is exactly U_f . We can use the same proof technique in Lemma 5 to show that $\overline{\text{ATI}}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$ can also be approximated by polynomial classical samples from D up to inverse polynomial precision and thus we can also apply hybrid arguments between $\overline{\text{ATI}}_{\mathcal{P}, D_0, \gamma}^{\epsilon, \delta}$ and $\overline{\text{ATI}}_{\mathcal{P}, D_1, \gamma}^{\epsilon, \delta}$ for computationally indistinguishable distributions D_0 and D_1 . Formally,

Lemma 6. *Let $\mathcal{H}_R, \mathcal{H}_{\text{Aux}}$ be Hilbert spaces. Let \mathcal{P} be a collection of projective measurements indexed by some set \mathcal{I} . Suppose \mathcal{P} can be implemented by a quantum circuit of size $|\mathcal{P}|$. Let D_0, D_1 be two efficiently sampleable distributions over \mathcal{I} . For any state $\rho \in \mathcal{D}(\mathcal{H}_R)$, denote $(b, \rho') \leftarrow \overline{\text{ATI}}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}(\rho)$ be the procedure that runs $\overline{\text{ATI}}_{\mathcal{P}, D, \gamma}^{\epsilon, \delta}$ on state ρ along with enough fresh ancillas initialized to $|0\rangle$, and gets an output b and the post-measurement state ρ' . For any polynomial μ , any quantum state ρ and any (possibly quantum) predicate $h : \{0, 1\} \times \mathcal{D}(\mathcal{H}_{\text{Aux}}) \rightarrow \{0, 1\}$ with circuit size $|h|$, if*

$$\left| \Pr[h(b, \rho') = 1 | (b, \rho') \leftarrow \overline{\text{ATI}}_{\mathcal{P}, D_0, \gamma}^{\epsilon, \delta}(\rho)] - \Pr[h(b, \rho') = 1 | (b, \rho') \leftarrow \overline{\text{ATI}}_{\mathcal{P}, D_1, \gamma}^{\epsilon, \delta}(\rho)] \right| \geq \frac{1}{\mu(\lambda)}.$$

Then there exists a quantum circuit \mathcal{C} of size $\text{poly}(\lambda, 1/\epsilon, \log(1/\delta), \mu, |\mathcal{P}|, |h|)$ (which only uses the quantum circuits to implement \mathcal{P} , h and to sample D_0, D_1 as a black box) such that

$$\left| \Pr[\mathcal{C}(\rho, x) = b : b \xleftarrow[\mathcal{P}, D, \gamma]{\epsilon, \delta} \{0, 1\}] - \frac{1}{2} \right| \geq \frac{1}{(\mu(\lambda))^3 \cdot \text{poly}(\lambda, 1/\epsilon, \log(1/\delta))}$$

which is an inverse polynomial if μ is a polynomial.

Proof. We omit the proof as it's almost the same as the proof of Lemma 5.

6 Almost Perfect Extraction of Preimages

In this section, we show how to extract a short preimage of \mathbf{y} with overwhelming probability, given a good (quantum) distinguisher between the distribution of a ciphertext of message μ and a uniform distribution. Our main contribution is an extraction algorithm that is guaranteed to work with *overwhelming probability*, in contrast to the extraction algorithm in [APV23] that only works with probability inversely proportional to the field size.

Since a general quantum distinguisher can be a superposition of a good distinguisher and a useless distinguisher, we use (Approximate) Threshold Implementation to (approximately) test whether a given quantum distinguisher is good before we apply the extraction algorithm. We need the following notations before we formally define what is a good quantum distinguisher.

Threshold Implementation on a Quantum Distinguisher. For a quantum algorithm \mathcal{A} with auxiliary quantum state ρ , let projective measurements $\{\mathcal{P}_x^{\mathcal{A}} = (P_x^{\mathcal{A}}, Q_x^{\mathcal{A}})\}$ correspond to running \mathcal{A} on x and the auxiliary state ρ . Suppose that $P_x^{\mathcal{A}}$ corresponds to outcome 1 and $Q_x^{\mathcal{A}}$ corresponds to outcome 0.

For two distributions D_0 and D_1 , denote (D_0, D_1) to be the distribution of (b, x) where $b \stackrel{\$}{\leftarrow} \{0, 1\}$ and $x \sim D_b$. We say that (\mathcal{A}, ρ) is a γ -good quantum distinguisher for distributions D_0 and D_1 with support \mathcal{X} if and only if ρ passes the projector $\text{TI}_{1/2+\gamma}(\mathcal{P}_{(D_0, D_1)}^{\mathcal{A}})$. Here, we abuse the notation to define the POVM $\mathcal{P}_{(D_0, D_1)}^{\mathcal{A}} = (P_{(D_0, D_1)}^{\mathcal{A}}, Q_{(D_0, D_1)}^{\mathcal{A}})$ ⁷ such that

$$P_{(D_0, D_1)}^{\mathcal{A}} = \frac{P_{D_1}^{\mathcal{A}} + Q_{D_0}^{\mathcal{A}}}{2} = \frac{\sum_{x \in \mathcal{X}} \Pr[x \leftarrow D_1] P_x^{\mathcal{A}} + \sum_{x \in \mathcal{X}} \Pr[x \leftarrow D_0] Q_x^{\mathcal{A}}}{2},$$

$$Q_{(D_0, D_1)}^{\mathcal{A}} = I - P_{(D_0, D_1)}^{\mathcal{A}}.$$

In other words, $\mathcal{P}_{(D_0, D_1)}^{\mathcal{A}} = (P_{(D_0, D_1)}^{\mathcal{A}}, Q_{(D_0, D_1)}^{\mathcal{A}})$ is the POVM measurement (where $P_{(D_0, D_1)}^{\mathcal{A}}$ corresponds to output 1 and $Q_{(D_0, D_1)}^{\mathcal{A}}$ corresponds to output 0) that on any input quantum state ρ ,

- Sample $(b, x) \sim (D_0, D_1)$.
- Feed x and the input quantum state ρ into \mathcal{A} , which outputs a guess b' .
- Output 1 if $b' = b$; 0 otherwise.

We denote the approximate version of $\text{TI}_{1/2+\gamma}(\mathcal{P}_{(D_0, D_1)}^{\mathcal{A}})$ as $\text{ATI}_{\mathcal{P}^{\mathcal{A}}, (D_0, D_1), 1/2+\gamma}^{\epsilon, \delta}$. Roughly speaking, $\text{ATI}_{\mathcal{P}^{\mathcal{A}}, (D_0, D_1), 1/2+\gamma}^{\epsilon, \delta}$ can efficiently estimate whether the algorithm \mathcal{A} , along with the input quantum state as auxiliary, can distinguish D_0 and D_1 with advantage at least γ . We denote the projective version of $\text{ATI}_{\mathcal{P}^{\mathcal{A}}, (D_0, D_1), 1/2+\gamma}^{\epsilon, \delta}$ as $\overline{\text{ATI}}_{\mathcal{P}^{\mathcal{A}}, (D_0, D_1), 1/2+\gamma}^{\epsilon, \delta}$.

⁷ $\mathcal{P}_{(D_0, D_1)}^{\mathcal{A}}$ is actually a mixture of projective measurements for the distribution (D_0, D_1) and a collection of binary outcome projective measurements $\mathcal{P}_{b, x} = (Q_x^{\mathcal{A}}, P_x^{\mathcal{A}})$ if $b = 0$ and $\mathcal{P}_{b, x} = (P_x^{\mathcal{A}}, Q_x^{\mathcal{A}})$ if $b = 1$.

Some Important Distributions. The threshold implementation will be used to test whether a quantum distinguisher works well on the following distributions. The prime modulus q , the noise ratios $\alpha, \beta \in (0, 1)$ and $n, m \in N$ are all fixed parameters that will be soon clear from the context. For matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and vectors $\mathbf{y} \in \mathbb{Z}_q^n, \mathbf{x} \in \mathbb{Z}_q^m$,

- Denote $D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}}$ to be the distribution of $(\mathbf{A}, \mathbf{y}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{s}^\top \mathbf{y} + e')$ where $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n, \mathbf{e} \sim D_{\mathbb{Z}_q^m, \alpha q}$ and $e' \sim D_{\mathbb{Z}_q, \beta q}$.
- Denote $D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}$ to be the distribution of $(\mathbf{A}, \mathbf{y}, \mathbf{u}^\top, u')$ where $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and $u' \xleftarrow{\$} \mathbb{Z}_q$.
- Denote $\bar{D}_{\text{lwe}}^{\mathbf{A}, \mathbf{x}}$ to be the distribution of $(\mathbf{A}, \mathbf{Ax}, \mathbf{u}^\top, \mathbf{u}^\top \mathbf{x} + e')$ where $e' \sim D_{\mathbb{Z}_q, \beta q}$ and $\mathbf{u}^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top$ for $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n, \mathbf{e} \sim D_{\mathbb{Z}_q^m, \alpha q}$.
- Denote $D_{\text{gl}}^{\mathbf{A}, \mathbf{x}}$ to be the distribution of $(\mathbf{A}, \mathbf{Ax}, \mathbf{u}^\top, \mathbf{u}^\top \mathbf{x} + e')$ where $e' \sim D_{\mathbb{Z}_q, \beta q}$ and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$.

For each of the above distribution D , we denote $D(i, g_i)$ to be the distribution of $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 + c \cdot \hat{\mathbf{i}}^\top, v_4 + c \cdot g_i)$ where $\hat{\mathbf{i}}$ is the unit vector with its i^{th} coordinate being 1, $c \xleftarrow{\$} \mathbb{Z}_q$, and $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, v_4) \sim D$. It is easy to generate a sample from $D(i, g_i)$ given i, g_i and a sample from D . Thus if we can efficiently distinguish between $D_0(i, g_i)$ and $D_1(i, g_i)$, then on input (i, g_i) , we can efficiently distinguish between D_0 and D_1 .

The adversary can be described as an unitary \mathcal{A} acting on a Hilbert space $\mathcal{H} = \mathcal{H}_{A_\lambda} \otimes \mathcal{H}_{B_\lambda} = \mathcal{H}_{R_\lambda} \otimes \mathcal{H}_{\text{Aux}_\lambda}$.

- H_{A_λ} stores the secret key state given by the challenger.
- H_{B_λ} is initialized to a quantum advice.
- H_{R_λ} stores the state returned to the challenger.
- H_{Aux_λ} is kept by the adversary.

We will omit λ when it is clear from the context. We show the following result.

Theorem 6 (Almost Optimal 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 $\lambda \in \mathbb{N}$ such that $m \leq \text{poly}(\lambda)$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and let $\alpha, \beta \in (0, 1)$ be noise ratios with $\beta/\alpha = 2^{o(n)}$, $2^{-o(n)} \leq \alpha\sigma \leq \text{negl}(\lambda)$ and $\alpha\sigma/\beta \leq \text{negl}(\lambda)$. 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 and polynomial-sized advice states $\nu_\lambda \in \mathcal{D}(\mathcal{H}_{B_\lambda})$ which are independent of \mathbf{A} and \mathbf{y} .*

Assume the decisional $\text{LWE}_{n, q, \alpha q}^m$ cannot be solved by a quantum algorithm running in time $\text{poly}(\lambda, \sigma)$ with distinguishing advantage $1/\text{poly}(\lambda, \sigma)$. If there exist functions $\varepsilon(\lambda) = 1/\text{poly}(\lambda)$, $\gamma(\lambda) = 1/\text{poly}(\lambda)$, $\delta(\lambda) = 2^{-\Theta(\lambda)}$ and a QPT distinguisher \mathcal{D} such that (Fig. 2)

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

SearchToDecisionExpt^{A,D}(1^λ):

1. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$.
2. Generate $(|\psi_y\rangle, y) \leftarrow \text{GenGauss}(\mathbf{A}, \sigma)$.
3. Generate $\rho_{R, \text{Aux}} \leftarrow \mathcal{A}_{\lambda, \mathbf{A}, y}(|\psi_y\rangle\langle\psi_y| \otimes \nu_\lambda)$.
4. Let two-outcome projective measurements $\{\mathcal{P}_x^D\}$ correspond to running \mathcal{D} on samples x and the auxiliary state in Aux. We approximately test whether $(\mathcal{D}, \rho_{\text{Aux}})$ is a γ -good quantum distinguisher between the distributions $D_{\text{lwe}}^{\mathbf{A}, y}$ and $D_{\text{unif}}^{\mathbf{A}, y}$ by running $\text{ATI}_{\mathcal{P}^D, (D_{\text{lwe}}^{\mathbf{A}, y}, D_{\text{unif}}^{\mathbf{A}, y}), 1/2+\gamma}^{\gamma/6, \delta}$ on ρ_{Aux} and output the result.

Fig. 2. The experiment SearchToDecisionExpt^{A,D}(1^λ).

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

$$\Pr \left[\mathbf{x} \in \Lambda_q^y(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2}) : \begin{array}{l} \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m} \\ (|\psi_y\rangle, y) \leftarrow \text{GenGauss}(\mathbf{A}, \sigma) \\ \rho_{R, \text{Aux}} \leftarrow \mathcal{A}_{\lambda, \mathbf{A}, y}(|\psi_y\rangle\langle\psi_y| \otimes \nu_\lambda) \\ 1 \leftarrow \text{ATI}_{\mathcal{P}^D, (D_{\text{lwe}}^{\mathbf{A}, y}, D_{\text{unif}}^{\mathbf{A}, y}), 1/2+\gamma}^{\gamma/6, \delta}(\rho_{\text{Aux}}) \\ \mathbf{x} \leftarrow \mathcal{E}(\mathbf{A}, y, \text{Aux}) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

6.1 Construction of the Extractor

In the subsection, we formally define our quantum extractor \mathcal{E} . \mathcal{E} takes \mathbf{A}, y and the quantum state in AUX as input, and does the following:

$\mathcal{E}(\mathbf{A}, y, \text{Aux})$:

1. Set $\mathbf{x} = \mathbf{0}$, $\epsilon = \gamma/6$, $\gamma' = \gamma - 3\epsilon = \gamma/2$.
2. For each $i = 1, 2, \dots, m$:
 - For each $g_i \in [-\sigma\sqrt{m/2}, \dots, \sigma\sqrt{m/2}]$:
 - i. Let AUX store the current state of the quantum distinguisher.
 - ii. Run $\overline{\text{ATI}}_{\mathcal{P}^D, (D_{\text{lwe}}^{\mathbf{A}, y}(i, g_i), D_{\text{unif}}^{\mathbf{A}, y}), 1/2+\gamma'}^{\epsilon, \delta}$ on the residual state in register AUX along with enough fresh ancillas initialized to $|0\rangle$.
 - iii. If it outputs 1, set $\mathbf{x}_i = g_i$, and move on to the next guess.
 - iv. If it outputs 0, move on to the next guess.
 3. Output \mathbf{x} .

Fig. 3. The quantum extractor $\mathcal{E}(\mathbf{A}, y, \text{Aux})$.

By construction, the extractor runs in time $\text{poly}(\lambda, \sigma, m, 1/\epsilon, \log \frac{1}{\delta}) = \text{poly}(\lambda, \sigma, 1/\gamma)$.

6.2 Analysis of the Extractor

Before we analyze the success probability of our extractor, we make crucial observations on the distributions $D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}}$, $D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}$, $\bar{D}_{\text{lwe}}^{\mathbf{A}, \mathbf{x}}$ and $D_{\text{gl}}^{\mathbf{A}, \mathbf{x}}$.

Lemma 7. *For any $\mathbf{x} \in \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2})$, the statistical distance between $\bar{D}_{\text{lwe}}^{\mathbf{A}, \mathbf{x}}$ and $D_{\text{lwe}}^{\mathbf{A}, \mathbf{Ax}}$ is at most $\text{negl}(\lambda)$.*

Proof. By noise smudging (Lemma 2), the statistical distance between the distribution $D_{\mathbb{Z}_q, \beta q}$ and the distribution of $\mathbf{e}^\top \mathbf{x} + e'$ where $e' \sim D_{\mathbb{Z}_q, \beta q}$ and $|\mathbf{e}^\top \mathbf{x}| \leq \alpha q \sigma m$ is at most $\alpha \sigma m / \beta$. Notice that for any $\mathbf{x} \in \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2})$, when \mathbf{e} is sampled from $D_{\mathbb{Z}_q^m, \alpha q}$, $|\mathbf{e}^\top \mathbf{x}| \geq \alpha q \sigma m$ with probability at most $2^{-\Omega(\lambda)}$ (from Banaszczyk's tail bound [Ban93]).

Thus the statistical distance between $\bar{D}_{\text{lwe}}^{\mathbf{A}, \mathbf{x}}$ and $D_{\text{lwe}}^{\mathbf{A}, \mathbf{Ax}}$ is at most $\alpha \sigma m / \beta + 2^{-\Omega(\lambda)}$, which by our choice of parameters, is at most $\text{negl}(\lambda)$.

Lemma 8. *For integer $i \in [m]$ and $g_i = \mathbf{x}_i$, $D_{\text{gl}}^{\mathbf{A}, \mathbf{x}}(i, g_i) = D_{\text{gl}}^{\mathbf{A}, \mathbf{x}}$.*

For integer $i \in [m]$ and $g_i \neq \mathbf{x}_i$, $D_{\text{gl}}^{\mathbf{A}, \mathbf{x}}(i, g_i) = D_{\text{unif}}^{\mathbf{A}, \mathbf{Ax}}$.

Proof. This follows directly from the definition, so we omit the proof.

Now we are ready to prove Theorem 6.

Proof. To prove Theorem 6, it suffices to prove that

$$\Pr \left[1 \leftarrow \text{Game}_0^{\mathcal{A}, \mathcal{D}}(1^\lambda) \right] \leq \text{negl}(\lambda)$$

where $\text{Game}_0^{\mathcal{A}, \mathcal{D}}$ is shown in Fig. 4.




Fig. 4. The game $\text{Game}_0^{\mathcal{A}, \mathcal{D}}(1^\lambda)$.

Let's consider the following sequence of hybrid distributions.

H_0 : This is the same as the game $\text{Game}_0^{\mathcal{A}, \mathcal{D}}(1^\lambda)$ defined in Fig. 4.

H_1 : This is the following distribution:

1. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$.
2. **Sample a Gaussian vector $\mathbf{x}' \sim D_{\mathbb{Z}_q^m, \sigma/\sqrt{2}}$ and let $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}' \bmod q$.**
3. Generate $\rho_{R, \text{Aux}} \leftarrow \mathcal{A}_{\lambda, \mathbf{A}, \mathbf{y}}(|\mathbf{x}'\rangle\langle\mathbf{x}'| \otimes \nu_\lambda)$.
4. Compute $b \leftarrow \overline{\text{ATI}}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{we}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma}^{\epsilon, \delta}(\rho_{\text{Aux}})$. Abort if $b = 0$.
5. Compute $\mathbf{x} \leftarrow \mathcal{E}(\mathbf{A}, \mathbf{y}, \text{Aux})$.
6. Output 1 if $\mathbf{x} \notin \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2})$; Otherwise, output 0.

H_2 : This is the following distribution:

1. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$.
2. Sample a Gaussian vector $\mathbf{x}' \sim D_{\mathbb{Z}_q^m, \sigma/\sqrt{2}}$ and let $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}' \bmod q$.
3. Generate $\rho_{R, \text{Aux}} \leftarrow \mathcal{A}_{\lambda, \mathbf{A}, \mathbf{y}}(|\mathbf{x}'\rangle\langle\mathbf{x}'| \otimes \nu_\lambda)$.
4. Compute $b \leftarrow \overline{\text{ATI}}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma}^{\epsilon, \delta}(\rho_{\text{Aux}})$. Abort if $b = 0$.
5. Compute $\mathbf{x} \leftarrow \mathcal{E}(\mathbf{A}, \mathbf{y}, \text{Aux})$.
6. Output 1 if $\mathbf{x} \notin \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2})$; Otherwise, output 0.

$H_{3,k}$: This is the following distribution which replaces $\overline{\text{ATI}}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{we}}^{\mathbf{A}, \mathbf{y}}(i, g_i), D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma}^{\epsilon, \delta}$ in \mathcal{E} with $\overline{\text{ATI}}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}(i, g_i), D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma}^{\epsilon, \delta}$ one by one (recall the description of \mathcal{E} defined in Fig. 3).

1. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$.
2. Sample a Gaussian vector $\mathbf{x}' \sim D_{\mathbb{Z}_q^m, \sigma/\sqrt{2}}$ and let $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}' \bmod q$.
3. Generate $\rho_{R, \text{Aux}} \leftarrow \mathcal{A}_{\lambda, \mathbf{A}, \mathbf{y}}(|\mathbf{x}'\rangle\langle\mathbf{x}'| \otimes \nu_\lambda)$.
4. Compute $b \leftarrow \overline{\text{ATI}}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma}^{\epsilon, \delta}(\rho_{\text{Aux}})$. Abort if $b = 0$.
5. Set $\mathbf{x} = \mathbf{0}$, $\epsilon = \gamma/6$, $\gamma' = \gamma - 3\epsilon = \gamma/2$, $\mathbf{t} = \mathbf{0}$.
6. For each $i = 1, 2, \dots, m$:
 - For each $g_i \in [-\sigma\sqrt{m/2}, \dots, \sigma\sqrt{m/2}]$:
 - i. Let AUX store the current state of the quantum distinguisher. $\mathbf{t} \leftarrow \mathbf{t} + 1$.
 - ii. If $\mathbf{t} \leq k$, run $\overline{\text{ATI}}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}(i, g_i), D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma}^{\epsilon, \delta}$ on the residual state in register AUX along with enough fresh ancillas initialized to $|0\rangle$. Otherwise, run $\overline{\text{ATI}}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{we}}^{\mathbf{A}, \mathbf{y}}(i, g_i), D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma}^{\epsilon, \delta}$ on the residual state in register AUX along with enough fresh ancillas initialized to $|0\rangle$.
 - iii. If it outputs 1, set $\mathbf{x}_i = g_i$, and move on to the next guess.
 - iv. If it outputs 0, move on to the next guess.
7. Output 1 if $\mathbf{x} \notin \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2})$; Otherwise, output 0.

We now show the following:

Lemma 9. *Assuming the quantum hardness of $\text{LWE}_{n,q,\alpha q}^m$, the hybrids H_0 and H_1 are computationally indistinguishable,*

$$H_0 \approx_c H_1.$$

Proof. This follows directly from the Gaussian-collapsing property (Theorem 2) for quantum distinguishers with auxiliary states.

Lemma 10. *Assuming the quantum hardness of $\text{LWE}_{n,q,\alpha q}^m$, the hybrids H_1 and H_2 are computationally indistinguishable,*

$$H_1 \approx_c H_2.$$

Proof. We prove the claim by contradiction.

Suppose H_1 and H_2 can be distinguished by a QPT algorithm \mathcal{B} with advantage $1/\lambda^c$ for a constant $c > 0$ and infinitely many λ . Fix one such λ .

By standard averaging argument, for at least $\frac{1}{2\lambda^c}$ fraction of $(\mathbf{A}, \mathbf{x}')$ sampled according to $\mathbf{A} \xleftarrow{s} \mathbb{Z}_q^{n \times m}$ and $\mathbf{x}' \sim D_{\mathbb{Z}_q^m, \sigma/\sqrt{2}}$, \mathcal{B} can distinguish the result of running step 3–6 of H_1 on $(\mathbf{A}, \mathbf{x}')$, and the result of running step 3–6 of H_2 on $(\mathbf{A}, \mathbf{x}')$ with advantage at least $\frac{1}{2\lambda^c}$. Let's call those $(\mathbf{A}, \mathbf{x}')$ good. Then from Lemma 5, there exists a quantum circuit \mathcal{C} of size $\text{poly}(\lambda, 1/\epsilon, \log(1/\delta))$ such that for each good $(\mathbf{A}, \mathbf{x}')$, $\mathcal{C}(\rho_{\text{Aux}}, \mathbf{A}, \mathbf{x}', \cdot)$ can distinguish samples from $(D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}, D_{\text{unif}}^{\mathbf{A}, \mathbf{x}'})$ and samples from $(D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}})$ with advantage at least $\frac{1}{\text{poly}(\lambda, 1/\epsilon, \log(1/\delta))}$.

As we can sample $D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}$ by ourselves and $D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}} \approx_s \bar{D}_{\text{lwe}}^{\mathbf{A}, \mathbf{x}'}$ (from Lemma 7 and the choice of parameters), there exists a polynomial size quantum circuit \mathcal{C}' such that for each good $(\mathbf{A}, \mathbf{x}')$, $\mathcal{C}'(\rho_{\text{Aux}}, \mathbf{A}, \mathbf{x}', \cdot)$ can distinguish samples from $D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}$ and $\bar{D}_{\text{lwe}}^{\mathbf{A}, \mathbf{x}'}$ with advantage at least $1/\lambda^d$ for some constant $d > 0$.

Recall that the only difference in $D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}$ and $\bar{D}_{\text{lwe}}^{\mathbf{A}, \mathbf{x}'}$ is whether \mathbf{u} is sampled according to LWE or sampled uniformly. Now let's show how to leverage the fact to break $\text{LWE}_{n,q,\alpha q}^m$ using this \mathcal{C}' (Algorithm 1). Notice that for all the good $(\mathbf{A}, \mathbf{x}')$, line 3 passes with noticeable probability (by averaging arguments over the eigenspaces) and the residual state after running ATI and obtaining outcome 1 is still a good distinguisher (by Lemma 4). So Algorithm 1 breaks decisional $\text{LWE}_{n,q,\alpha q}^m$ efficiently if Lemma 10 doesn't hold.

This ends our proof of the claim.

Lemma 11. *Assume that the decisional $\text{LWE}_{n,q,\alpha q}^m$ cannot be solved by a quantum algorithm running in time $\text{poly}(\lambda, \sigma)$ with distinguishing advantage $1/\text{poly}(\lambda, \sigma)$.*

The probability that hybrid $H_{3,k}$ outputs 1 and the probability that hybrid $H_{3,k+1}$ outputs 1 are $\text{negl}(\lambda)/\sigma$ close. Formally, for $0 \leq k \leq \sqrt{2}\sigma m^{3/2} - 1$,

$$|\Pr[H_{3,k+1} = 1] - \Pr[H_{3,k} = 1]| \leq \text{negl}(\lambda)/\sigma$$

Proof. The proof is the same with the proof of Lemma 10 except that we apply Lemma 6 instead of Lemma 5 and that we set the parameter μ in Lemma 6 as $1/\text{poly}(\lambda, \sigma)$ instead of $1/\text{poly}(\lambda)$. We omit the proof details.

Algorithm 1: An algorithm to break decisional $\text{LWE}_{n,q,\alpha q}^m$ if Lemma 10 doesn't hold

Input : Matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and vector $\mathbf{u} \in \mathbb{Z}_q^n$ (and quantum advice ν_λ).
Output: 0 or 1 (guess whether \mathbf{u} is sampled from uniform or according to $\text{LWE}_{n,q,\alpha q}^m$)

- 1 Sample a vector $\mathbf{x}' \sim D_{\mathbb{Z}_q^m, \sigma/\sqrt{2}}$ and let $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}' \bmod q$.
- 2 Generate $\rho_{R, \text{Aux}} \leftarrow \mathcal{A}_{\lambda, \mathbf{A}, \mathbf{y}}(|\mathbf{x}'\rangle\langle\mathbf{x}'| \otimes \nu_\lambda)$.
- 3 Test whether $\mathcal{C}'(\rho_{\text{Aux}}, \mathbf{A}, \mathbf{x}', \cdot)$ can be used to distinguish samples from $D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}$ and samples from $\bar{D}_{\text{LWE}}^{\mathbf{A}, \mathbf{x}'}$ with advantage at least $1/\lambda^d$ by running $\text{ATI}^{1/\lambda^{d+1}, \delta}$ on it with threshold $1/2 + \frac{1}{4\lambda^d}$. If the ATI outputs 0 (it is not a good distinguisher), output a random guess and abort.
- 4 Denote the residual state (if not abort) in register AUX as ρ'_{Aux} .
- 5 Sample $e' \sim D_{\mathbb{Z}_q, \beta q}$. Let $\mathbf{v} := (\mathbf{A}, \mathbf{A}\mathbf{x}', \mathbf{u}^\top, \mathbf{u}^\top \mathbf{x}' + e')$.
- 6 Run $\mathcal{C}'(\rho'_{\text{Aux}}, \mathbf{A}, \mathbf{x}', \mathbf{v})$ and output the result.

Lemma 12. $\mathsf{H}_{3, \sqrt{2}\sigma m^{3/2}}$ outputs 1 with negligible probability.

Proof. We first define $\text{Game}_1^{\mathcal{A}, \mathcal{D}}(1^\lambda)$ in Fig. 5. It is the same as $\mathsf{H}_{3, \sqrt{2}\sigma m^{3/2}}$ except that it will output 1 if $\mathbf{x} \neq \mathbf{x}'$ (which is implied by $\mathbf{x} \notin \mathcal{A}_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2})$), so to prove Lemma 12, it suffices to prove that $\text{Game}_1^{\mathcal{A}, \mathcal{D}}(1^\lambda)$ outputs 1 with negligible probability.




Fig. 5. The game $\text{Game}_1^{\mathcal{A}, \mathcal{D}}(1^\lambda)$.

Notice that in step 6, we apply a sequence of projective measurements $\overline{\text{ATI}}$ and set each coordinate of \mathbf{x}' based on the measurement outcomes. By Quantum Union Bound (Lemma 1), $\Pr[1 \leftarrow \text{Game}_1^{\mathcal{A}, \mathcal{D}}(1^\lambda)]$ can be bounded by a union of events that for \mathbf{x}' sampled according to $D_{\mathbb{Z}_q^m, \sigma/\sqrt{2}}$, $\text{SubGame}^{\mathcal{A}, \mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')$ outputs 1:

$$\begin{aligned} & \Pr[1 \leftarrow \text{Game}_1^{\mathcal{A}, \mathcal{D}}(1^\lambda)] \\ & \leq 4 \sum_{i=1}^m \sum_{g_i=-\sigma\sqrt{m/2}}^{\sigma\sqrt{m/2}} \Pr[1 \leftarrow \text{SubGame}^{\mathcal{A}, \mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}') : \mathbf{x}' \sim D_{\mathbb{Z}_q^m, \sigma/\sqrt{2}}] \end{aligned}$$

where $\text{SubGame}^{\mathcal{A}, \mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')$ is defined in Fig. 6.




Fig. 6. The game $\text{SubGame}^{\mathcal{A}, \mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')$.

Now let's show for any fixed i, g_i, \mathbf{x}' ,

$$\Pr[1 \leftarrow \text{SubGame}^{\mathcal{A}, \mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')] \leq \text{negl}(\lambda)/\sigma \quad (1)$$

Case 1: $g_i = \mathbf{x}'_i$ Consider the residual state ρ'_{Aux} of running step 3 and obtaining $b = 1$. From Lemma 4, running $\text{ATI}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma'}^{\epsilon, \delta}$ on ρ'_{Aux} , we will obtain 1 with probability at least $1 - 3\delta$.

From Lemma 8, when $g_i = \mathbf{x}'_i$, $D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}(i, g_i) = D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}$. Thus the output distribution of running $\text{ATI}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma'}^{\epsilon, \delta}$ on ρ'_{Aux} is exactly the same as that of running $\overline{\text{ATI}}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}(i, g_i), D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma'}^{\epsilon, \delta}$ on ρ'_{Aux} . Therefore,

$$\Pr[1 \leftarrow \text{SubGame}^{\mathcal{A}, \mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')] \leq 3\delta \leq \text{negl}(\lambda)/\sigma.$$

Case 2: $g_i \neq \mathbf{x}'_i$ Again consider the residual state ρ'_{Aux} of running step 3 and obtaining $b = 1$. From Lemma 8, when $g_i \neq \mathbf{x}'_i$, $D_{\text{gl}}^{\mathbf{A}, \mathbf{x}'}(i, g_i) = D_{\text{unif}}^{\mathbf{A}, \mathbf{A}\mathbf{x}'} = D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}$. Thus when running $\text{ATI}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma'}$ on ρ'_{Aux} , we will obtain 1 with probability exactly $\Pr[1 \leftarrow \text{SubGame}^{\mathcal{A}, \mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')]$.

As $\mathcal{P}_{(D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}})}^{\mathcal{D}}$ only has eigenvalue $1/2 < 1/2 + \gamma' - \epsilon$ (any distinguisher cannot do better than outputting a random guess when facing $D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}$ and $D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}$), from Lemma 4, running $\text{ATI}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2+\gamma'}$ on any state, we cannot get 1 with probability greater than δ , which implies that

$$\Pr[1 \leftarrow \text{SubGame}^{\mathcal{A}, \mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')] \leq \delta \leq \text{negl}(\lambda)/\sigma.$$

Summing up Eq. 1 and averaging over \mathbf{x}' , we can get that

$$\begin{aligned} & \Pr[1 \leftarrow \text{Game}_1^{\mathcal{A}, \mathcal{D}}(1^\lambda)] \\ & \leq 4 \sum_{i=1}^m \sum_{g_i=-\sigma\sqrt{m/2}}^{\sigma\sqrt{m/2}} \Pr[1 \leftarrow \text{SubGame}^{\mathcal{A}, \mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}') : \mathbf{x}' \sim D_{\mathbb{Z}_q^m, \sigma/\sqrt{2}}] \\ & \leq \text{negl}(\lambda), \end{aligned}$$

which ends the proof.

Recall that H_0 is the same as the game $\text{Game}_0^{\mathcal{A}, \mathcal{D}}(1^\lambda)$, Theorem 6 follows directly from Lemma 9, Lemma 10, Lemma 11 over $\Theta(\sigma m^{3/2})$ pairs of consecutive hybrids, Lemma 12 and the observation that $\mathsf{H}_2 = \mathsf{H}_{3,0}$.

7 Proof of Theorem 5

We prove by contradiction. Let \mathcal{A} be the QPT adversary and without loss of generality, we assume that the adversary submits $\mu = 0$ and assume that


$$\Pr[1 \leftarrow \text{Expt}^{\Sigma, \mathcal{A}}(1^\lambda, 1)] - \Pr[1 \leftarrow \text{Expt}^{\Sigma, \mathcal{A}}(1^\lambda, 0)] = \epsilon(\lambda),$$

where $\epsilon(\lambda)$ is inverse polynomial, $\text{Expt}^{\Sigma, \mathcal{A}}(1^\lambda, b)$ is defined as Fig. 1 and $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$.

We decompose the adversary into two QPT algorithms \mathcal{A}, \mathcal{D} where given input state $|\psi_{\mathbf{y}}\rangle$, \mathcal{A} generates the state $\rho_{\text{R}, \text{Aux}}$. After returning system R to the challenger, \mathcal{D} takes ρ_{Aux} and responds to the challenge. Then \mathcal{A}, \mathcal{D} satisfy

$$\Pr[1 \leftarrow \text{SecurityExpt}^{\mathcal{A}, \mathcal{D}}(1^\lambda, 1)] - \Pr[1 \leftarrow \text{SecurityExpt}^{\mathcal{A}, \mathcal{D}}(1^\lambda, 0)] = \epsilon(\lambda)$$

where $\text{SecurityExpt}^{\mathcal{A}, \mathcal{D}}$ is the experiment shown in Fig. 7, because the inefficient revocation implements $\text{Revoke}(\text{MSK}, \text{PK}, \sigma)$.

**Fig. 7.** The experiment SecurityExpt^{A,D} (1^λ, b).


Lemma 13. For adversary \mathcal{A}, \mathcal{D} that satisfy

$$\Pr \left[1 \leftarrow \text{SecurityExpt}^{\mathcal{A}, \mathcal{D}}(1^\lambda, 1) \right] - \Pr \left[1 \leftarrow \text{SecurityExpt}^{\mathcal{A}, \mathcal{D}}(1^\lambda, 0) \right] = \epsilon(\lambda),$$

they also satisfy

$$\Pr \left[1 \leftarrow \text{ATISecurityExpt}^{\mathcal{A}, \mathcal{D}, \gamma}(1^\lambda) \right] \geq \frac{\epsilon(\lambda)}{4} - \text{negl}.$$

for $\gamma = \frac{3\epsilon}{14}$ where ATISecurityExpt^{A,D} is shown in Fig. 8.

**Fig. 8.** The experiment ATISecurityExpt^{A,D,ε} (1^λ, b).

Proof. Suppose that revocation succeeds with probability p . The residual state ρ_{Aux} satisfies

$$\mathbb{E} \left[\text{Tr} \left[\mathcal{P}_{(D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}})}^{\mathcal{D}} \rho_{\text{Aux}} \right] \mid \text{Revocation succeeds on R} \right] \geq \frac{1}{2} + \frac{\epsilon}{2p}.$$

By averaging argument and the definition of threshold implementation Theorem 3,

$$\mathbb{E} \left[\text{Tr} \left[\text{TI}_{\frac{1}{2} + \frac{\epsilon}{4}} \left(\mathcal{P}_{(D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}})}^{\mathcal{D}} \right) \rho_{\text{AUX}} \right] \middle| \text{Revocation succeeds on R} \right] \geq \frac{\epsilon}{4p}.$$

By Lemma 4, if we set $\delta = 2^{-\Theta(\lambda)}$ we have,

$$\begin{aligned} & \Pr \left[\text{ATI}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), \frac{1}{2} + \gamma}^{\frac{\gamma}{6}, \delta} (\rho_{\text{AUX}}) = 1 \middle| \text{Revocation succeeds on R} \right] \\ &= \mathbb{E} \left[\text{Tr} \left[\text{ATI}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), \frac{1}{2} + \gamma}^{\frac{\gamma}{6}, \delta} \rho_{\text{AUX}} \right] \middle| \text{Revocation succeeds on R} \right] \\ &\geq \mathbb{E} \left[\text{Tr} \left[\text{TI}_{\frac{1}{2} + \frac{\epsilon}{4}} \left(\mathcal{P}_{(D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}})}^{\mathcal{D}} \right) \rho_{\text{AUX}} \right] \middle| \text{Revocation succeeds on R} \right] - \delta \\ &\geq \frac{\epsilon}{4p} - \text{negl}. \end{aligned}$$

Using the above lemma we can construct Algorithm 2 for solving $\text{SIS}_{n, q, \sigma\sqrt{2m}}^m$ problem using the adversary \mathcal{A}, \mathcal{D} . As for our choice of parameters, the hardness of $\text{LWE}_{n, q, \alpha q}^m$ implies the hardness of $\text{SIS}_{n, q, \sigma\sqrt{2m}}^m$, Theorem 5 follows directly from the correctness of Algorithm 2, which we show in the following claim.

Algorithm 2: 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 $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \text{GenGauss}(\mathbf{A}, \sigma)$ with

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

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

- 2 Run \mathcal{A} to generate a bipartite state $\rho_{R, \text{AUX}}$ in systems $\mathcal{H}_R \otimes \mathcal{H}_{\text{AUX}}$ with $\mathcal{H}_R = \mathcal{H}_q^m$.
- 3 Run $\text{ATI}_{\mathcal{P}^{\mathcal{D}}, (D_{\text{lwe}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2 + \gamma}^{\gamma/6, \delta}$ on system AUX , abort if the output is 0.
- 4 Run the extractor $\mathcal{E}(\mathbf{A}, \mathbf{y}, \text{AUX})$ from Theorem 6, and let $\mathbf{x}_1 \in \mathbb{Z}_q^n$ denote the outcome.
- 5 Measure system R in the computational basis, and let $\mathbf{x}_0 \in \mathbb{Z}_q^n$ denote the outcome.
- 6 Output the vector $\mathbf{x} = \mathbf{x}_1 - \mathbf{x}_0$.

Claim. Algorithm 2 solves $\text{SIS}_{n, q, \sigma\sqrt{2m}}^m$ with inverse polynomial probability when \mathcal{A}, \mathcal{D} is a successful adversary.

Proof. Suppose \mathcal{A}, \mathcal{D} is a successful adversary. To show that Algorithm 2 can obtain a short non-zero solution \mathbf{x} we prove the following two statements:

SimultExtractionExpt^{A, D}(1^λ):

1. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$.
2. Generate $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \text{GenGauss}(\mathbf{A}, \sigma)$.
3. Generate $\rho_{R, \text{Aux}} \leftarrow \mathcal{A}_{\lambda, \mathbf{A}, \mathbf{y}}(|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}| \otimes \nu_{\lambda})$.
4. Run $\text{ATI}_{\mathcal{P}^D, (D_{\text{low}}^{\mathbf{A}, \mathbf{y}}, D_{\text{unif}}^{\mathbf{A}, \mathbf{y}}), 1/2 + \gamma}^{\gamma/6, \delta}$ on system Aux , abort if the output is 0.
5. Run the extractor $\mathcal{E}(\mathbf{A}, \mathbf{y}, \text{AUX})$ from [Theorem 6](#), and let $\mathbf{x}_1 \in \mathbb{Z}_q^n$ denote the outcome. Abort if $\mathbf{x}_1 \notin \Lambda_q^{\mathbf{y}}(\mathbf{A}) \cap \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2})$.
6. Apply inefficient revocation on system R , abort if it fails; Otherwise, output 1.

Fig. 9. The experiment $\text{SimultExtractionExpt}^{\mathcal{A}, \mathcal{D}}(1^{\lambda})$.

- The probability that on system AUX the extractor \mathcal{E} extracts a short preimage \mathbf{x}_1 of \mathbf{y} and revocation succeeds on R is inverse polynomial

$$\Pr \left[\text{SimultExtractionExpt}^{\mathcal{A}, \mathcal{D}}(1^{\lambda}) = 1 \right] = \frac{1}{\text{poly}(\lambda)}.$$

where $\text{SimulExtractionExpt}$ is defined as [Fig. 9](#).

- Suppose that revocation succeeds with probability $\varepsilon(\lambda)$ conditioned on the extraction being successful. Then instead of running revocation on R , if we measure register R in computational basis and obtain result \mathbf{x}_0 , the probability that \mathbf{x}_0 is a short preimage of \mathbf{y} that is different from \mathbf{x}_1 is $\varepsilon(\lambda) - \text{negl}(\lambda)$ conditioned on the extraction being successful.

If both statements are true, by basic probability arguments we prove the claim.

The first statement follows from [Lemma 13](#) and [Theorem 6](#). Let GoodDecryptor denote the event that we pass the ATI test on step 4. Let RevocationSuc denote the event that the inefficient revocation succeeds on system R on step 6. Let ExtractionSuc denote the event that \mathbf{x}_1 is a short preimage of \mathbf{y} on step 5. Since step 4–5 and step 6 commute, by [Lemma 13](#),

$$\Pr [\text{RevocationSuc} \wedge \text{GoodDecryptor}] = \frac{1}{\text{poly}(\lambda)}.$$

By [Theorem 6](#),

$$\Pr [\text{ExtractionSuc} \mid \text{GoodDecryptor}] \geq 1 - \text{negl}(\lambda).$$

By basic probability calculation,

$$\Pr [\text{RevocationSuc} \wedge \text{ExtractionSuc}] = \frac{1}{\text{poly}(\lambda)}.$$

Now we prove the second statement. We show that given a specific short preimage \mathbf{x}_1 of \mathbf{y} and a state ρ_R such that revocation succeeds on R with probability $\varepsilon(\lambda)$, if we measure R under computational basis, we obtain a short preimage \mathbf{x}_0 of \mathbf{y} that is different from \mathbf{x}_1 with probability $\varepsilon(\lambda) - \text{negl}(\lambda)$. Define the set of short preimages $\mathcal{S} = \{\mathbf{x} | \mathbf{Ax} = \mathbf{y}, \|\mathbf{x}\| \leq \sigma\sqrt{\frac{m}{2}}\}$ and

$$|\psi'_y\rangle = \left(\sum_{\mathbf{x} \in \mathcal{S}} \rho_{\frac{\sigma}{\sqrt{2}}}(\mathbf{x}) \right)^{-\frac{1}{2}} \sum_{\mathbf{x} \in \mathcal{S}} \rho_\sigma(\mathbf{x}) |\mathbf{x}\rangle$$

be a ‘truncated’ Gaussian coset state. Consider the following projectors

- $\Pi_0 = \sum_{\mathbf{x} \in \mathcal{S}, \mathbf{x} \neq \mathbf{x}_1} |\mathbf{x}\rangle\langle\mathbf{x}|$ is a projector that projects onto all short preimages we want.
- $\Pi_1 = |\psi'_y\rangle\langle\psi'_y|$ is the approximate revocation projector. The trace distance between Π_1 and the actual revocation projector $|\psi_y\rangle\langle\psi_y|$ is negligible by Banaszczky’s tail bound [Ban93].

Suppose that \mathbf{A} is a full-rank matrix, if $\text{Tr}[|\psi_y\rangle\langle\psi_y|\rho_R] = \varepsilon$ we have

$$\begin{aligned} \text{Tr}[\Pi_0 \rho_R] &\geq \text{Tr}[\Pi_1 \Pi_0 \rho_R] \\ &\geq \text{Tr}[\Pi_1 \rho_R] - \text{negl}(\lambda) \\ &\geq \text{Tr}[|\psi_y\rangle\langle\psi_y|\rho_R] - \text{negl}(\lambda) \\ &= \varepsilon - \text{negl}(\lambda). \end{aligned}$$

where the second inequality follows from Lemma 3. Note that \mathbf{A} is full-rank with $1 - \text{negl}(\lambda)$ probability. Combine all arguments above, we proof this claim.

8 Applications

Combining our result with [APV23], we obtain constructions for

- Public-Key Encryption with Classical Key Revocation.
- Key-Revocable Fully Homomorphic Encryption.
- Revocable Pseudorandom Functions.

8.1 Public-Key Encryption with Classical Key Revocation

A public-key encryption with classical key-revocation is a public-key encryption such that whenever we want to perform key revocation:

- The lessee runs `Delete` on its quantum secret key ρ_{SK} and produce a classical certificate π .
- The lessor runs `Revoke` on input π and output `Valid` if it is a valid certificate.

The security of such scheme captures the idea that if an adversary produces a certificate π that passes the revocation then the remaining adversary cannot distinguish between a ciphertext of chosen message from a random ciphertext. In [APV23], they built a public-key encryption with classical key-revocation assuming the security of key-revocable Dual-Regev encryption. Combine with our result, we obtain the following theorem.

Theorem 7. *Assuming the polynomial hardness of LWE with sub-exponential modulus. The scheme $\text{CRevDual} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Delete}, \text{Revoke})$ (Construction 2, [APV23]) is a secure public-key encryption with classical key-revocation (Definition 7.1, 7.2, [APV23]).*

8.2 Key-Revocable Fully Homomorphic Encryption

A key-revocable fully homomorphic encryption is a fully homomorphic encryption with quantum key revocation just like the key-revocable Dual-Regev Encryption. In [APV23], they built a key-revocable fully homomorphic encryption assuming the security of key-revocable Dual-Regev encryption. Meanwhile, this construction can be adapted to feature classical revocation via techniques used in public-key encryption with classical key-revocation mentioned above. Combine with our result, we obtain the following theorem.

Theorem 8. *Assuming the polynomial hardness of LWE and SIS with sub-exponential modulus. The scheme $\text{RevDualGSW} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval}, \text{Revoke})$ (Construction 3, [APV23]) is a secure key-revocable fully homomorphic encryption (Definition 5.3, [APV23]). Meanwhile, this construction can be adapted to feature classical revocation via (Construction 2, [APV23]).*

8.3 Revocable Pseudorandom Functions

A key-revocable (or simply, revocable) pseudorandom function is a weak pseudorandom function with its evaluation key revocable. The μ -security of such scheme captures the idea that if the revocation succeeds, the remaining adversary cannot distinguish between μ images $y_1 = \text{PRF}(x_1), y_2 = \text{PRF}(x_2), \dots, y_\mu = \text{PRF}(y_\mu)$ from μ random preimages x_1, x_2, \dots, x_μ and uniform random values y_1, y_2, \dots, y_μ . Meanwhile, this construction can also be adapted to feature classical revocation. Combine with our result, we obtain the following theorem.

Theorem 9. *Assuming the polynomial hardness of LWE and SIS with sub-exponential modulus. The scheme $(\text{Gen}, \text{PRF}, \text{Eval}, \text{Revoke})$ (Construction 5, [APV23]) is a poly-secure revocable PRF scheme (Definition 9.2, 9.3, [APV23]). Meanwhile, this construction can be adapted to feature classical revocation via (Construction 2, [APV23]).*

Acknowledgments. We thank Jiahui Liu for patiently answering our questions about her work [CGJL23] and thank anonymous reviewers for their helpful comments.

Prabhanjan Ananth is supported by the National Science Foundation under Grant No. 2329938. This work was done while Zihan Hu was an intern at Shanghai Qi Zhi Institute.

References

- [Aar09] Aaronson, S.: Quantum copy-protection and quantum money. In: 2009 24th Annual IEEE Conference on Computational Complexity, pp. 229–242. IEEE (2009)
- [Aar16] Aaronson, S.: The Complexity of Quantum States and Transformations: From Quantum Money to Black Holes (2016). [arXiv:1607.05256](https://arxiv.org/abs/1607.05256) [quant-ph]
- [AGKZ20] Amos, R., Georgiou, M., Kiayias, A., Zhandry, M.: One-shot signatures and applications to hybrid quantum/classical authentication. In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, pp. 255–268 (2020)
- [AHH24] Ananth, P., Hu, Z., Huang, Z.: Quantum key- revocable dual-regev encryption, revisited. Cryptology ePrint Archive, Paper 2024/738 (2024). <https://eprint.iacr.org/2024/738>
- [AKN+23] Agrawal, S., Kitagawa, F., Nishimaki, R., Yamada, S., Yamakawa, T.: Public key encryption with secure key leasing. arXiv preprint [arXiv:2302.11663](https://arxiv.org/abs/2302.11663) (2023)
- [ALL+21] Aaronson, S., Liu, J., Liu, Q., Zhandry, M., Zhang, R.: New approaches for quantum copy-protection. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12825, pp. 526–555. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84242-0_19
- [APV23] Ananth, P., Poremba, A., Vaikuntanathan, V.: Revocable cryptography from learning with errors. In: Rothblum, G., Wee, H. (eds) TTCC 2023, Part IV. LNCS, vol. 14372, pp. 93–122. Springer, Cham (2023). ISBN 978-3-031-48623-4. https://doi.org/10.1007/978-3-031-48624-1_4
- [Ban93] Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen* **296**(4), 625–636 (1993). <http://eudml.org/doc/165105>
- [BCM+21] Brakerski, Z., Christiano, P., Mahadev, U., Vazirani, U., Vidick, T.: A cryptographic test of quantumness and certifiable randomness from a single quantum device (2021). [arXiv:1804.00640](https://arxiv.org/abs/1804.00640) [quant-ph]
- [BDGM20] Brakerski, Z., Döttling, N., Garg, S., Malavolta, G.: Factoring and pairings are not necessary for IO: circularsecure LWE suffices. Cryptology ePrint Archive (2020)
- [BGG+14] Boneh, D., et al.: Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 533–556. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_30
- [BI20] Broadbent, A., Islam, R.: Quantum encryption with certified deletion. In: Pass, R., Pietrzak, K. (eds.) TCC 2020. LNCS, vol. 12552, pp. 92–122. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64381-2_4
- [BK22] Bartusek, J., Khurana, D.: Cryptography with certified deletion (2022). <https://doi.org/10.48550/ARXIV.2207.01754>, <https://arxiv.org/abs/2207.01754>
- [BL20] Broadbent, A., Lord, S.: Uncloneable quantum encryption via oracles. In: Flammia, S.T. (ed.) 15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020), vol. 158. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl– Leibniz-Zentrum für Informatik, pp. 4:1–4:22 (2020). <https://doi.org/10.4230/LIPIcs.TQC.2020.4>

- [CGJL23] Chardouvelis, O., Goyal, V., Jain, A., Liu, J.: Quantum key leasing for PKE and FHE with a classical lessor. arXiv preprint [arXiv:2310.14328](https://arxiv.org/abs/2310.14328) (2023)
- [DGT+10] Dodis, Y., Goldwasser, S., Kalai, Y.T., Peikert, C., Vaikuntanathan, V.: Public-key encryption schemes with auxiliary inputs. In: Micciancio, D. (eds.) TCC 2010. LNCS, vol. 5978, pp. 361–381. Springer, Cham (2010). ISBN 978-3-642-11799-2. https://doi.org/10.1007/978-3-642-11799-2_22
- [Die82] Dieks, D.G.B.J.: Communication by EPR devices. *Phys. Lett. A* **92**(6), 271–272 (1982)
- [Gao15] Gao, J.: Quantum union bounds for sequential projective measurements. *Phys. Rev. A* **92**(5), 052331 (2015)
- [GPV08] Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, pp. 197–206 (2008)
- [HMNY21] Hiroka, T., Morimae, T., Nishimaki, R., Yamakawa, T.: Quantum encryption with certified deletion, revisited: public key, attribute-based, and classical communication (2021). [arXiv: 2105.05393](https://arxiv.org/abs/2105.05393) [quant-ph]
- [LMZ23] Liu, J., Montgomery, H., Zhandry, M.: Another round of breaking and making quantum money: how to not build it from lattices, and more. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023. LNCS, vol. 14004, pp. 611–638. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30545-0_21
- [Mah18] Mahadev, U.: Classical homomorphic encryption for quantum circuits. In: Thorup, M. (ed.) 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, 7–9 October 2018, pp. 332–338. IEEE Computer Society (2018). <https://doi.org/10.1109/FOCS.2018.00039>
- [MW05] Marriott, C., Watrous, J.: Quantum Arthur-Merlin Games (2005). [arXiv: cs/0506068](https://arxiv.org/abs/cs/0506068) [cs.CC]
- [Por22] Poremba, A.: Quantum proofs of deletion for learning with errors (2022). <https://doi.org/10.48550/ARXIV.2203.01610>
- [Qua20] Quach, W.: UC-secure OT from LWE, revisited. In: Galdi, C., Kolesnikov, V. (eds.) SCN 2020. LNCS, vol. 12238, pp. 192–211. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-57990-6_10
- [WZ82] Wootters, W.K., Zurek, W.H.: A single quantum cannot be cloned. *Nature* **299**(5886), 802–803 (1982)
- [Zha20] Zhandry, M.: Schrödinger’s pirate: how to trace a quantum decoder. Cryptology ePrint Archive, Paper 2020/1191. <https://eprint.iacr.org/2020/1191.2020>