Simple and Efficient SKL-IBE with Classical Revocation from LWE

Ho Nguyen Pham[1], Duong Hieu Phan[1], Quoc-Huy Vu[2], and Weiqiang Wen[1]

[1] LTCI, Telecom Paris, Institut Polytechnique de Paris, France
{ho.pham,hieu.phan,weiqiang.wen}@telecom-paris.fr
[2] De Vinci Higher Education, De Vinci Research Center, Paris, France
quoc.huy.vu@ens.fr

Abstract. Secure key leasing (SKL) is a quantum cryptographic primitive that enables the leasing of decryption keys to delegated users with the guarantee that, once revoked, the lessees irreversibly lose decryption capability. A key feature that makes SKL practically relevant is *classical revocation*: the ability to revoke keys at any time and from anywhere, without relying on a quantum channel.
In this work, we revisit SKL schemes for *public-key encryption* (PKE) and *identity-based encryption* (IBE), and present a new approach for concrete efficiency under the standard Learning With Errors (LWE) assumption. First, we refine the security analysis of the Dual-Regev SKL-PKE scheme with classical revocation from [Ananth, Poremba, and Vaikuntanathan, TCC 2023; Ananth, Hu, and Huang, TCC 2024], establishing security under the polynomial hardness of LWE with polynomial modulus. Together with the resource efficiency of the Dual-Regev-based construction, our analysis shows that this approach yields the most quantum-efficient known SKL-PKE scheme. Second, we present a simple and efficient construction of *selectively secure SKL-IBE with classical revocation from standard LWE*. Our approach directly extends the Dual-Regev SKL-PKE within the IBE framework of [Agrawal, Boneh, and Boyen, Eurocrypt 2010], avoiding garbled circuits and obfuscation-based assumptions and achieving improved concrete efficiency over prior generic approaches.

Keywords: Secure key leasing · Identity-based encryption · Learning with errors · Dual-Regev encryption

1 Introduction

Secure Key Leasing (SKL), introduced in [KN22, AKN+23, APV23], is a quantum cryptographic primitive that leverages the no-cloning theorem [WZ82] to enable the leasing and revocation of secret keys. In such a scheme, a lessor (Alice) leases a decryption key to a lessee (Bob), who may use it until Alice requests revocation. Upon revocation, Bob must return a verifiable certificate; SKL security requires that once this certificate is accepted, Bob irreversibly loses the ability to decrypt. This primitive is highly desirable for scenarios where

temporary access to sensitive information is needed, such as granting employees access during contract terms or providing subscription-based access to digital content. Crucially, SKL provides an effective solution for verifiable revocation of decryption rights *without requiring re-encryption or key updates.*

Identity-Based Encryption (IBE) was proposed [Sha84] as an advanced public-key encryption primitive, where users' public keys can be derived from arbitrary strings, such as email addresses or usernames. In this setting, the ability to efficiently revoke malicious users is desirable and has motivated the study of revocable IBE (RIBE) [BGK08]. However, RIBE schemes by definition require key updates and can only protect ciphertexts encrypted with the updated keys. By contrast, SKL-IBE leverages quantum information to support verifiable revocation of decryption capabilities without such restrictions.

Classical revocation. Revocation in SKL can be *classical* or *quantum*, depending on whether the returned certificate is a classical string or a quantum state. As argued in prior works on unclonable cryptography [Shm22, CHV23], even in a world where quantum computers are widely accessible, classical communication channels remain far more practical and cost-effective than quantum communications. As in our motivating scenarios, while key distribution can occur when parties are physically close, forcing them to be present at specific locations with quantum computers during revocation seems highly impractical (e.g., when the lessor is traveling and has no access to quantum resources). Classical certificates are compatible with efficient and simultaneous (potentially public) verification, whereas quantum certificates typically require sequential verification and dedicated quantum-channel verifiers, incurring considerable overhead. For these reasons, achieving classical revocation is both practically desirable and theoretically important.

Dual-Regev Encryption with SKL. The [APV23] introduces the notion of key-revocable PKE (equivalent to SKL-PKE) and gives a simple construction based on the Dual-Regev PKE. Their proof relies on LWE with subexponential hardness and subexponential modulus with additional conjectures; the [AHH24] removes the conjectures and improves the analysis to polynomial hardness, but still requires subexponential modulus.

On the other hand, a key advantage of the Dual-Regev SKL-PKE is that verifying the revocation for a Gaussian coset state does not require parallel repetition to amplify soundness. By contrast, other known approaches to achieve SKL-PKE rely on this soundness amplification technique to achieve comparable security guarantees, including constructions based on BB84 states [AKN+23, KMY25] and on the noisy trapdoor claw-free functions (NTCF) [PWYZ24, CGJL25], which force the lessee to maintain at least λ quantum keys[3] where λ is the security parameter. Since leased keys must often be obtained and stored

[3] In [KMY25], instantiation with λ-qubit BB84 states is necessary for 2^λ-level security.

long before use, the required *long-term quantum memory* is therefore a central concrete metric.

Existing approaches for SKL-IBE. Prior work gives generic SKL constructions for ABE/FE [AKN+23, BGK+24, KNP25b], which imply IBE, but these routes often trade concrete efficiency or rely on stronger assumptions. In particular, the [AKN+23] obtains SKL-ABE (and hence SKL-IBE) by combining any SKL-PKE with a classical ABE/IBE scheme via garbling: ciphertexts include a garbled SKL-PKE encryption circuit whose labels are encrypted under IBE. As a result, the scheme requires 2ℓ parallel IBE instances, where $\ell = |\mathsf{pk}|$ is the bit-length of the SKL-PKE public key. While ℓ is asymptotically polynomial in the security parameter, the value of ℓ can range from thousands to hundreds of thousands when instantiated with post-quantum assumptions. Another generic approach in [KNP25b] uses MI-ABE and compute-and-compare obfuscation to obtain classical revocation and collusion resistance, but the required MI-ABE variant is only known from indistinguishability obfuscation.

Our results. In light of the above limitations, our contributions are as follows.

1. We refine the security analysis of the Dual-Regev SKL-PKE scheme with classical revocation from [APV23, AHH24], establishing security under the polynomial hardness of LWE with *polynomial modulus*. Specifically, we employ a gentle noise-smudging analysis, inspired by [BD20, PWYZ24], to handle LWE errors. Together with the resource efficiency of the Dual-Regev-based construction, our analysis shows that this approach yields the most efficient known SKL-PKE scheme in terms of quantum resources (Table 1).
2. We present a simple and efficient construction of *selectively secure* SKL-IBE[4] with classical revocation under the standard LWE assumption, achieving improved *concrete efficiency* over prior generic approaches. Compared with the generic construction of [AKN+23], our approach avoids garbling entirely, thereby eliminating both the 2ℓ-fold IBE overhead and the $O(s)$ garbling cost, where s denotes the size of the SKL-PKE encryption circuit.

For comparison, our SKL-PKE scheme preserves the $O(\lambda \log^2 \lambda)$-qubit complexity of the Dual-Regev-based scheme while targeting the same 2^λ-level security as [AKN+23, KMY25, CGJL25, PWYZ24]. By contrast, obtaining this qubit complexity from the analyses of [APV23, AHH24] requires a quasi-polynomial modulus, which results in a weaker concrete security guarantee.

[4] As in [AKN+23], our SKL-IBE security experiment considers a single key query for the challenge identity. See Remark 1 for details.

[5] Although [CGJL25, PWYZ24] achieve zero quantum communication complexity, they require substantially more quantum memory than ours. In some architectures (e.g., linear optics), quantum memory is effectively implemented via communication channels and can be even more costly in practice.

Table 1: Comparison of SKL-PKE. Quantum resources account for both the quantum communication complexity and the quantum memory required to store the quantum keys.[5]

Construction	Key Distribution	Revocation	Assumption	Quantum Resources (qubits)
[AKN+23]	Quantum	Quantum	PKE	$\Omega(\lambda^2)$
[KMY25]	Quantum	Classical	PKE	$\Omega(\lambda^2)$
[APV23, AHH24]	Quantum	Classical	LWE	$\omega(\lambda \log^2 \lambda)$[1]
[CGJL25, PWYZ24]	Classical	Classical	LWE	$O(\lambda^2 \log \lambda)$
Ours	Quantum	Classical	LWE	$O(\lambda \log^2 \lambda)$

(1) This bound requires an aggressive parameter choice with a quasi-polynomial modulus.

1.1 Technical Overview

Starting point. We begin by recalling the Dual-Regev SKL-PKE construction of [APV23], henceforth APV, given by efficient algorithms $(\mathcal{KG}, \mathsf{Enc}, \mathcal{Dec}, \mathcal{Del}, \mathsf{Verif})$. The quantum $\mathcal{KG}$ algorithm generates a classical Dual-Regev public key $(\mathbf{A}, \mathbf{y})$ and a quantum secret key - a Pauli-$\mathbf{Z}$-twirled Gaussian coset state:

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

where $\rho_\sigma(\mathbf{x}) = \exp\left(-\pi \|x\|^2 / \sigma^2\right)$ is the Gaussian measure for $\sigma > 0$, $\mathbf{v} \leftarrow\!\$\ \{0,1\}^m$ is a random m-bit string, and $\mathbf{Z}_q$ is the q-ary Pauli-$\mathbf{Z}$ operator. In particular, we prepare a Gaussian superposition over $\mathbf{x} \in \mathbb{Z}_q^m$, compute $\mathbf{A}\mathbf{x}$ into an auxiliary register and measure it to collapse onto a random coset $\{\mathbf{x} \in \mathbb{Z}_q^m : \mathbf{A}\mathbf{x} = \mathbf{y}\}$. Encryption follows Dual-Regev [GPV08], while decryption is performed coherently using $\left|\psi^{\mathbf{v}}_{\mathbf{A},\mathbf{y}}\right\rangle$ (the $\mathbf{Z}_q$ twist affects only the phase and does not change correctness). Revocation applies a Fourier transform on the quantum key and measures the state, resulting in a shifted LWE sample $\mathbf{w}^\intercal = \hat{\mathbf{s}}^\intercal \mathbf{A} + \hat{\mathbf{e}}^\intercal + \left\lfloor \frac{q}{\nu} \right\rfloor \cdot \mathbf{v}^\intercal$. The trapdoor $\mathbf{T}_\mathbf{A}$ enables efficient recovery of $\mathbf{v}$ from $\mathbf{w}$ for verification.

Dual-Regev SKL-PKE with polynomial modulus. We now explain where the subexponential modulus requirement enters the APV analysis. The main source is the noise smudging technique used to argue statistical indistinguishability between $\left(\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}_0^\intercal, \mathbf{s}^\intercal \mathbf{A}\mathbf{x} + e'\right)$ and $\left(\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}_0^\intercal, \mathbf{s}^\intercal \mathbf{A}\mathbf{x} + \mathbf{e}_0^\intercal \mathbf{x} + e'\right)$. Their argument sets the bound of e' to exceed that of $\mathbf{e}_0^\intercal \mathbf{x}$ by a superpolynomial factor, resulting in a subexponential modulus. Our main observation is that the smudging noise e' need not be sampled independently of $\mathbf{x}$. Then, inspired by [BD20], we decompose e' into $\mathbf{e}_0^\intercal \mathbf{x} + e''$, where the norm bound of e'' is only polynomially larger than that of $\mathbf{e}_0^\intercal \mathbf{x}$. This idea was also used by [PWYZ24] to obtain polynomial modulus for SKL in a related setting, but our technical context differs from both works. A second source is the collapsing analysis underlying the Gaussian coset state's uncertainty principle: while earlier proofs assumed LWE with subexponential modulus [LZ19, Por23], we rely on [LMZ23] to obtain the same guarantee under polynomial modulus.

SKL-IBE from LWE. We next consider a SKL-IBE scheme with classical revocation given by efficient algorithms $(\mathsf{Setup}, \mathcal{KG}, \mathsf{Enc}, \mathcal{Dec}, \mathcal{Del}, \mathsf{Verif})$. In particular, Setup outputs $(\mathsf{mpk}, \mathsf{msk})$, $\mathsf{Enc}(\mathsf{mpk}, \mathsf{id}, \mu)$ encrypts to an identity id, and $\mathcal{KG}(\mathsf{msk}, \mathsf{id})$ outputs a quantum decryption key ρ_{id} used by $\mathcal{Dec}$. Revocation is captured by a deletion algorithm $\mathcal{Del}(\rho_{\mathsf{id}})$, which outputs a *classical* certificate cert, and a classical verification algorithm $\mathsf{Verif}(\mathsf{msk}, \mathsf{id}, \mathsf{cert})$ that checks its validity. We further require *stateless* verification, meaning that Verif maintains no evolving private state across key generations.

From APV, we follow the standard lattice-based IBE template [ABB10, CHKP10]. For each identity id, we define $\mathbf{A}_{\mathsf{id}} = [\mathbf{A} \mid \bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B}]$, where $\mathbf{A}, \bar{\mathbf{A}}, \mathbf{B} \in \mathbb{Z}_q^{n \times m}$ and $\mathsf{H}(\cdot)$ maps identities to matrices in $\mathbb{Z}_q^{n \times n}$. We set $\mathsf{mpk} = (\mathbf{A}, \bar{\mathbf{A}}, \mathbf{B}, \mathbf{y})$ and encrypt as $(\mathbf{s}^\intercal \mathbf{A}_{\mathsf{id}} + \mathbf{e}^\intercal[\mathbf{I} \mid \mathbf{R}] + [\mathbf{0} \mid \mu])$ for some random $\mathbf{R} \in \{-1, 1\}^{m \times m}$. A natural approach would be generating a coset state $\left|\psi^{\mathbf{v}}_{\mathbf{A}_{\mathsf{id}}, \mathbf{y}}\right\rangle$ as in APV for the decryption key. However, in IBE the syndrome $\mathbf{y}$ must be fixed before key queries, so generating such states requires a quantum analogue of Gaussian preimage sampling using a trapdoor for $\mathbf{A}_{\mathsf{id}}$. We use the preimage-sampling procedure of APV as a subroutine and then measure the second half of the resulting state, resulting in a Gaussian coset state $\left|\psi_{\mathbf{A}, \hat{\mathbf{y}}}\right\rangle$ for $\mathbf{A}$ together with a classical outcome $\mathbf{u}$. We then apply a q-ary Pauli-$\mathbf{Z}$ twist and define the identity key as

$$\rho_{\mathsf{id}} = (\left|\psi^{\mathbf{v}}_{\mathbf{A}, \hat{\mathbf{y}}}\right\rangle, \hat{\mathbf{y}}, \mathbf{u}), \text{ where } \hat{\mathbf{y}} + (\bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B})\mathbf{u} = \mathbf{y}. \tag{1}$$

We will justify this form in the security proof discussion. For revocation, we follow APV, except that we sample $\mathbf{v} \leftarrow \mathsf{F}(\mathsf{k}, \mathsf{id} \| \hat{\mathbf{y}})$, where F is a pseudorandom function and k is the secret seed included in msk, yielding stateless verification.

Security proof. We prove selective security under the standard LWE assumption. In this setting, the adversary commits to a target identity id^* before seeing mpk and may query quantum decryption keys for any identity, *including* id^*. When the challenger requests revocation, the adversary must return a certificate for id^* and can only receive the challenge ciphertext under id^* if the certificate is accepted. Afterwards, the adversary may continue to request keys for any identity except id^*. Security requires that no QPT adversary distinguishes a real challenge ciphertext from random with non-negligible advantage over $1/2$.

Following APV, we prove security by reducing the IND-SKL adversary to the uncertainty principle for Gaussian coset states: given $\left|\psi^{\mathbf{v}}_{\mathbf{A}, \mathbf{y}}\right\rangle$ for uniformly random $\mathbf{A}$ with hidden phase $\mathbf{v}$, no adversaries $(\mathcal{B}_0, \mathcal{B}_1)$ satisfy (i) QPT algorithm $\mathcal{B}_0$ extracts a short $\mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{y}$ and (ii) computationally unbounded algorithm $\mathcal{B}_1$ recovers $\mathbf{v}$ from the internal state of $\mathcal{B}_0$. Our reduction must therefore simulate key queries without access to the trapdoor $\mathbf{T}_{\mathbf{A}}$. Since id^* is fixed, we program $\bar{\mathbf{A}} := \mathbf{A}\mathbf{R}^* - \mathsf{H}(\mathsf{id}^*)\mathbf{B}$ for some random $\mathbf{R}^* \in \{-1, 1\}^{m \times m}$. For any $\mathsf{id} \neq \mathsf{id}^*$, $\mathbf{A}_{\mathsf{id}} = [\mathbf{A} \mid \mathbf{A}\mathbf{R}^* + (\mathsf{H}(\mathsf{id}) - \mathsf{H}(\mathsf{id}^*))\mathbf{B}]$ and a trapdoor $\mathbf{T}_{\mathbf{A}_{\mathsf{id}}}$ can be extended from a trapdoor $\mathbf{T}_{\mathbf{B}}$ whenever $\mathsf{H}(\mathsf{id}) - \mathsf{H}(\mathsf{id}^*)$ is full rank. This programming technique, inspired by [ABB10], requires additional care for SKL. First, $\mathbf{A}^*_{\mathsf{id}} = [\mathbf{A} \mid \mathbf{A}\mathbf{R}^*]$ is not uniformly random from the reduction's view since it needs $\mathbf{R}^*$ for the

challenge ciphertext. As the uncertainty principle holds only for uniform $\mathbf{A}$, we need to formulate the leased key using coset states with respect to $\mathbf{A}$ as in Equation (1). Second, unlike classical IBE, we allow $\mathcal{A}$ to request the decryption key for id^*. However, since we program $\mathbf{A}_{\mathsf{id}^*} = [\mathbf{A} \mid \mathbf{AR}^*]$ to embed LWE pseudorandomness in the challenge ciphertext, our reduction has no trapdoor to answer this query! To handle this issue, we prepare ρ_{id^*} beforehand by generating $\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle$ and $\mathbf{u} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma/\sqrt{2}}$, then programming $\mathbf{y} = \hat{\mathbf{y}} + \mathbf{AR}^*\mathbf{u}$ in mpk. This allows us to sidestep the challenge of answering query for *decrypting* key, with the caveat that the proof handles only a single key query for id^*, while there is no limit on the number of key queries for $\mathsf{id} \neq \mathsf{id}^*$. We note that a new technique to handle the programming of $\mathbf{y}$ could address the multiple key queries challenge and may also be useful toward adaptive security without relying on complexity leveraging. We leave this as an interesting direction for future work.

We are now ready to complete the reduction. Given the challenge state $\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle$, $\mathcal{B}_0$ programs mpk and simulates the view of $\mathcal{A}$ until it receives a valid certificate $\mathbf{w}$ for ρ_{id^*}, then runs a QPT extractor $\mathcal{E}$ on the residual state of $\mathcal{A}$ to extract a short $\mathbf{x} : \mathbf{Ax} = \hat{\mathbf{y}}$. $\mathcal{B}_1$ can then inefficiently recover $\mathbf{v}$ from the certificate $\mathbf{w}_{\mathsf{id}^*}$, thus contradicting the uncertainty principle. At a high level, this $\mathcal{E}$ treats $\mathcal{A}$ as a decision oracle distinguishing the real and random challenge ciphertext distributions $\left(\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}_0^\intercal, (\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}_0^\intercal)\mathbf{R}^*, \mathbf{s}^\intercal\mathbf{y} + e'\right)$ and $\left(\mathbf{c}_0^\intercal, \mathbf{c}_1^\intercal, c_2'\right)$, where $\mathbf{y} = \hat{\mathbf{y}} + \mathbf{AR}^*\mathbf{u}$. An astute reader might notice that, if we can argue the indistinguishability between $\left(\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}_0^\intercal, (\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}_0^\intercal)\mathbf{R}, \mathbf{s}^\intercal\mathbf{y} + e'\right)$ and $\left(\mathbf{c}_0^\intercal, \mathbf{c}_1^\intercal, \mathbf{c}_0^\intercal\mathbf{x} + \mathbf{c}_1^\intercal\mathbf{u} + e'\right)$ with $\mathbf{c}_0 \leftarrow\!\!\$\ \mathbb{Z}_q^m$, extracting $\mathbf{x}$ from such distinguisher resembles the search-to-decision reduction for LWE introduced by Regev [Reg05] via coordinate-wise guessing. This technique was first used for SKL in [CGJL25] and underlies the extractor of [AHH24]. Our setting requires a separate analysis because the distinguisher operates on different distributions (see Section 4.3).

1.2 Related Work

Collusion Resistance (CR). In [AKN+23], the security against bounded CR of *decrypting keys* was considered and their transformation from 1-bounded to Q-bounded CR translates to our SKL-IBE. However, [KNP25b] is the first to address the stronger notion of unbounded CR. Recently, [KNP25a] targets the same notion but for primitives other than encryption. Achieving *unbounded CR SKL-IBE with classical revocation* from *standard assumptions* remains open.

Classical Communication. The line of work on SKL with completely classical communication [PWYZ24, CGJL25, KLYY25] relies on the adaptive hardcore bit property of the LWE-based NTCF to dequantize both key generation and revocation. This approach, originating from [BCM+18, Mah18], is the only known route to give completely classical communication for SKL to date.

SKL from ideal/module lattices. Similar to classical IBE, moving to ideal or module lattices may yield more compact and efficient schemes. However, to the

best of our knowledge, no SKL scheme from structured lattices has been proposed yet, despite existing results on NTCF from Ring-LWE [BKVV20, LG22].

2 Preliminaries

Notation. Let λ denote the security parameter. We write $\mathsf{negl}(\lambda)$ for any f with $f(\lambda) = \lambda^{-\omega(1)}$ and $\mathsf{poly}(\lambda)$ for any f with $f(\lambda) = \mathcal{O}(\lambda^c)$ for some $c > 0$. For $a, b \in \mathbb{R}, n \in \mathbb{N}$, let $[a, b] := \{x \in \mathbb{Z} \mid a \leq x \leq b\}$ and $[n] := [1, n]$. We write $a \leftarrow\!\!\$\, \mathcal{U}$ for uniform sampling from set $\mathcal{U}$ and $a \leftarrow \mathcal{A}$ for sampling from a probabilistic algorithm $\mathcal{A}$. By PPT (resp., QPT), we mean a polynomial-time non-uniform family of probabilistic (resp., quantum) circuits. We use bold font for (column) vectors and matrices, sans serif font for classical algorithms and variables, and script font for quantum algorithms. For two random variables X and Y, $X \approx_s Y$ and $X \approx_c Y$ denote statistical and computational indistinguishability.

2.1 Secure Key Leasing (SKL)

Definition 1 (SKL-IBE with Classical Revocation). *A secure key leasing identity-based encryption scheme with classical revocation is a tuple of algorithms* $\Sigma = (\mathsf{Setup}, \mathcal{KG}, \mathsf{Enc}, \mathcal{Dec}, \mathcal{Del}, \mathsf{Verif})$*, defined as follows with* $\{0,1\}^*$ *being the identity space:*

- $\mathsf{Setup}(1^\lambda) \to (\mathsf{mpk}, \mathsf{msk})$*: On input the security parameter* λ*, output a master public key* mpk *and a master secret key* msk*.*
- $\mathcal{KG}(\mathsf{msk}, \mathsf{id}) \to \rho_{\mathsf{id}}$*: On input the master secret key* msk *and an identity* id*, output a quantum state* ρ_{id}*.*
- $\mathsf{Enc}(\mathsf{mpk}, \mathsf{id}, \mu \in \{0,1\}) \to \mathsf{ct}$*: On input the master public key* mpk*, an identity* id*, and a message* μ*, output a ciphertext* ct*.*
- $\mathcal{Dec}(\rho_{\mathsf{id}}, \mathsf{ct}) \to \{\mu', \bot\}$*: On input a quantum state* ρ_{id} *and a ciphertext* ct*, output a message* μ' *if decryption is successful, otherwise output* $\bot$*.*
- $\mathcal{Del}(\rho_{\mathsf{id}}) \to \mathsf{cert}_{\mathsf{id}}$*: On input a quantum state* ρ_{id}*, output a certificate* $\mathsf{cert}_{\mathsf{id}}$*.*
- $\mathsf{Verif}(\mathsf{msk}, \mathsf{cert}_{\mathsf{id}}) \to \{\top, \bot\}$*: On input the master secret key* msk *and a certificate* $\mathsf{cert}_{\mathsf{id}}$*, output a bit indicating whether the certificate is valid or not.*

Decryption Correctness. For any $\mathsf{id} \in \{0,1\}^*$ *and* $\mu \in \{0,1\}$*, the following holds:*

$$\Pr\left[\mu \leftarrow \mathcal{Dec}(\rho_{\mathsf{id}}, \mathsf{ct}) : \begin{array}{l}(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda),\\ \rho_{\mathsf{id}} \leftarrow \mathcal{KG}(\mathsf{msk}, \mathsf{id}),\\ \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{mpk}, \mathsf{id}, \mu)\end{array}\right] \geq 1 - \mathsf{negl}(\lambda).$$

Verification Correctness. For any $\mathsf{id} \in \{0,1\}^*$*, the following holds:*

$$\Pr\left[\top \leftarrow \mathsf{Verif}(\mathsf{msk}, \mathsf{cert}) : \begin{array}{l}(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda),\\ \rho_{\mathsf{id}} \leftarrow \mathcal{KG}(\mathsf{msk}, \mathsf{id}),\\ \mathsf{cert} \leftarrow \mathcal{Del}(\rho_{\mathsf{id}}),\end{array}\right] \geq 1 - \mathsf{negl}(\lambda).$$

Definition 2 (Selective IND-SKL security of SKL-IBE). *We define the following security experiment* $\mathsf{Expt}^{\mathsf{SKL\text{-}IBE}}_{\mathcal{A}}(1^\lambda, b)$*:*

1. $\mathcal{A}$ *sends a challenge identity* id^* *to* $\mathcal{Ch}$*.*
2. $\mathcal{Ch}$ *runs* $(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda)$*, sends* mpk *to* $\mathcal{A}$*.*
3. $\mathcal{A}$ *makes key generation queries to* $\mathcal{O}_{\mathcal{KG}}$ *with at most 1 query for* $\mathsf{id} = \mathsf{id}^*$*.* $\mathcal{O}_{\mathcal{KG}}(\mathsf{id})$*: returns* $\rho_{\mathsf{id}} \leftarrow \mathcal{KG}(\mathsf{msk}, \mathsf{id})$*.*
4. $\mathcal{Ch}$ *sends* REVOKE *to* $\mathcal{A}$*, then* $\mathcal{A}$ *sends a classical certificate* $\mathsf{cert}_{\mathsf{id}^*}$ *and the plaintext* μ *to* $\mathcal{Ch}$*.*
5. $\mathcal{Ch}$ *runs* Verif *on the certificate and continues if* $\mathsf{Verif}(\mathsf{msk}, \mathsf{cert}_{\mathsf{id}^*}) \to \top$*, otherwise outputs* $\perp$ *and aborts.*
6. $\mathcal{Ch}$ *computes* $\mathsf{ct}_0 \leftarrow \mathsf{Enc}(\mathsf{mpk}, \mathsf{id}^*, \mu)$ *and samples* $\mathsf{ct}_1 \leftarrow\!\!\$\ \mathcal{C}$ *(*$\mathcal{C}$ *is the ciphertext space), then sends* ct_b *to* $\mathcal{A}$*.*
7. $\mathcal{A}$ *can continue to make queries to* $\mathcal{O}_{\mathcal{KG}}$*, except for* $\mathsf{id} = \mathsf{id}^*$*.*
8. $\mathcal{A}$ *outputs a guess* b' *and* $\mathcal{Ch}$ *outputs* b' *as the output of the experiment.*

A SKL-IBE scheme is selectively IND-SKL secure *if for every* QPT *adversary* $\mathcal{A}$*, the following holds:*

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

Remark 1. Our security experiment only allows a single key query for the challenge identity id^*, which is equivalent to the security notions considered in [AKN+23]. Allowing polynomial-bounded number of key queries for id^* could be achieved using the same transformation as in [AKN+23], but we do not consider it here for simplicity.

2.2 Quantum Background

We assume familiarity with quantum information and computation, and refer to [NC11] for the definition of basic concepts. We recall the following lemma which will be used in our security proof.

Lemma 1 (Quantum Union Bound, [Gao15]). *Let* $\rho \in \mathcal{D}(\mathcal{H})$ *be a state and let* $\Pi_1, \ldots, \Pi_n \geq 0$ *be a sequence of (orthogonal) projections acting on* $\mathcal{H}$*. Suppose that, for every* $i \in [n]$*, it holds that* $\mathrm{Tr}[\Pi_i\, \rho] = 1 - \varepsilon_i$*, for* $\varepsilon_i \in [0, 1]$*. Then, if we sequentially measure* ρ *with projective measurements* $\{\Pi_1, I - \Pi_1\}, \ldots, \{\Pi_n, I - \Pi_n\}$*, the probability that all* n *measurements succeed is at least*

$$\mathrm{Tr}\big[\Pi_n \cdots \Pi_1\, \rho\, \Pi_1 \cdots \Pi_n\big] \;\geq\; 1 \;-\; 4\sum_{i=1}^{n} \varepsilon_i.$$

2.3 Gaussian Coset States

A Gaussian coset state corresponding to a matrix $\mathbf{A}$ is defined as:

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

where $\mathbf{y} \in \mathbb{Z}_q^n$ is statistically close to uniform whenever $m \geq 2n \log q$ and $\sigma \geq \omega(\sqrt{\log m})$. We note that given as input a matrix $\mathbf{A}$ and $\sigma > 0$, there exists an efficient quantum algorithm $\mathit{GenGauss}(\mathbf{A}, \sigma)$ that generates $|\psi_{\mathbf{A},\mathbf{y}}\rangle$ and an efficient quantum algorithm $\mathit{SampGauss}(\mathbf{A}, \mathbf{T}_\mathbf{A}, \mathbf{y}, \sigma)$ that generates a state within negligible trace distance from $|\psi_{\mathbf{A},\mathbf{y}}\rangle$ [APV23, Por23]. Details on technical lemmas on Gaussian coset states are given in Appendix A.4.

3 Dual-Regev SKL-PKE with Polynomial Modulus

In this section, we revisit the Dual-Regev SKL-PKE scheme and refine the security analysis to obtain security under LWE with polynomial modulus. For completeness, we include the construction of [APV23].

Construction 1 (Dual-Regev SKL-PKE [APV23]). Let $q \geq 2$ be a prime, let $n, m, \nu \in \mathbb{N}$ and $\sigma, \chi, \chi' > 0$ be parameters. The construction for Dual-Regev SKL-PKE scheme $\mathsf{SKL\text{-}PKE} = (\mathcal{KG}, \mathsf{Enc}, \mathcal{Dec}, \mathcal{Del}, \mathsf{Verif})$ is defined as follows:

- $\mathcal{KG}(1^\lambda) \to (\mathsf{pk}, \mathsf{msk}, \rho_{\mathsf{sk}})$: Sample $(\mathbf{A}, \mathbf{T}_\mathbf{A}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$, $\mathbf{v} \leftarrow\!\!\$\ \{0,1\}^m$, and $(|\psi_{\mathbf{A},\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathit{GenGauss}(\mathbf{A}, \sigma)$. Output $\mathsf{pk} = (\mathbf{A}, \mathbf{y})$, $\mathsf{msk} = (\mathbf{T}_\mathbf{A}, \mathbf{v})$, and

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

- $\mathsf{Enc}(\mathsf{pk}, \mu \in \{0,1\}) \to \mathsf{ct}$: Sample $\mathbf{s} \leftarrow\!\!\$\ \mathbb{Z}_q^n$ and $\mathbf{e} \leftarrow D_{\mathbb{Z}^m,\chi}, e' \leftarrow D_{\mathbb{Z},\chi'}$. Output $\mathsf{ct} = (\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal, \mathbf{s}^\intercal \mathbf{y} + e' + \mu \cdot \lfloor q/2 \rfloor) \in \mathbb{Z}_q^{m+1}$.
- $\mathcal{Dec}(\rho_{\mathsf{sk}}, \mathsf{ct}) \to \{0,1\}$: Apply the unitary $U : |\mathbf{x}\rangle |\mathbf{0}\rangle \mapsto |\mathbf{x}\rangle |\mathsf{ct} \cdot (-\mathbf{x}, 1)\rangle$. Measure the second register in the computational basis. Output 0 if the measurement outcome is closer to 0 than to $\lfloor q/2 \rfloor$, otherwise output 1.
- $\mathcal{Del}(\rho) \to \mathsf{cert}$: Apply the m-qudit q-ary Fourier transform to the state ρ, then measure in the computational basis to obtain $\mathbf{w} \in \mathbb{Z}_q^m$. Output $\mathsf{cert} = \mathbf{w}$.
- $\mathsf{Verif}(\mathsf{msk}, \mathsf{cert}) \to \{\top, \bot\}$:
  - Parse $\mathsf{msk} = (\mathbf{T}_\mathbf{A} \in \mathbb{Z}_q^{m \times m}, \mathbf{v} \in \{0,1\}^m)$ and $\mathsf{cert} = \mathbf{w}$.
  - Compute $\mathbf{c}^\intercal = \mathbf{w}^\intercal \cdot \mathbf{T}_\mathbf{A} \pmod q$ and $\mathbf{d}^\intercal = \mathbf{c}^\intercal \cdot \mathbf{T}_\mathbf{A}^{-1}$, where $\mathbf{T}_\mathbf{A}^{-1} \in \mathbb{R}^{m \times m}$.
  - For $i \in [m]$, let $v'_i = 0$, if $d_i \in [-q/\sigma, q/\sigma]$, and let $v'_i = 1$, otherwise.
  - Output $\top$, if $\mathbf{v}' = \mathbf{v}$, and output $\bot$ otherwise.

Parameters. Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n, m \in \mathbb{N}$, q be a prime modulus with $m = \lceil 6n \log q \rceil$ and $\sqrt{8m} < \sigma < q/\sqrt{8m}$, and $\nu = 64m^2$, each parameterized by λ. Let $\chi, \chi' > 2\sqrt{n}$ such that $\chi'/\chi > \sigma\sqrt{2m}$.

Correctness. Decryption and verification correctness of the construction follow immediately from [APV23].

Theorem 1. *Assuming the quantum hardness of* LWE *with polynomial modulus, the scheme* $\mathsf{SKL\text{-}PKE} = (\mathcal{KG}, \mathsf{Enc}, \mathcal{Dec}, \mathcal{Del}, \mathsf{Verif})$ *in Construction 1 with specified parameters is IND-SKL secure.*

Since the proof of Theorem 1 is essentially the same as the proof of [AHH24, Theorem 5], we only provide details on new analysis for polynomial modulus and refer the reader to the original paper for the full proof.

First of all, we prove the security of the construction by showing that from a successful adversary for our SKL-PKE construction, we can win the uncertainty principle experiment in Theorem 2 (adapted from [APV23]).

Theorem 2 (Uncertainty principle for Gaussian coset states). *Let $n \in \mathbb{N}$ and q be a prime modulus with $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and $\nu > 0$. Assuming the quantum hardness of LWE with polynomial modulus, for any adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ consisting of a QPT algorithm $\mathcal{A}_0$ and an unbounded algorithm $\mathcal{A}_1$, we have:*

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

where $\mathsf{UncertaintyExpt}_{\mathcal{A}}(1^\lambda)$ is defined as follows:

1. *$\mathcal{Ch}$ samples a random $\mathbf{A} \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times m}$, a random $\mathbf{v} \leftarrow\!\!\$\ \{0,1\}^m$, generates $(|\psi_{\mathbf{A},\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathcal{GenGauss}(\mathbf{A}, \sigma)$, prepares $|\psi^{\mathbf{v}}_{\mathbf{A},\mathbf{y}}\rangle = \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} |\psi_{\mathbf{A},\mathbf{y}}\rangle$, and sends $(\mathbf{A}, \mathbf{y}, |\psi^{\mathbf{v}}_{\mathbf{A},\mathbf{y}}\rangle)$ to $\mathcal{A}_0$.*
2. *$\mathcal{A}_0$ generates a vector $\mathbf{x}_0 \in \mathbb{Z}_q^m$ and a polynomial-sized quantum state ρ_{aux}. Then, $\mathcal{A}_0$ sends $\mathbf{x}_0$ to $\mathcal{Ch}$ and forwards $(\mathbf{A}, \mathbf{y}, \rho_{\mathsf{aux}})$ to $\mathcal{A}_1$.*
3. *$\mathcal{A}_1$ generates a vector $\mathbf{v}' \in \{0,1\}^m$ on input $(\mathbf{A}, \mathbf{y}, \rho_{\mathsf{aux}})$ and sends it to $\mathcal{Ch}$.*
4. *$\mathcal{Ch}$ checks if $\mathbf{v}' = \mathbf{v}$ and whether $\mathbf{x}_0$ satisfies $\mathbf{A} \cdot \mathbf{x}_0 = \mathbf{y}$ and $\|\mathbf{x}_0\| \leq \sigma\sqrt{m/2}$. If true, $\mathcal{Ch}$ outputs 1, otherwise 0, as the outcome of the experiment.*

Proof. The proof is the same as [APV23, Theorem 7.7], except that we invoke Theorem 7, which relies on the quantum hardness of LWE with polynomial modulus, instead of the results relying on LWE with subexponential modulus. □

The proof of [AHH24, Theorem 5] continues to construct a QPT extractor that, conditioned on the adversary being a good distinguisher between real ciphertext and random distributions, can extract a short preimage $\mathbf{x}$ such that $\mathbf{Ax} = \mathbf{y}$ with overwhelming probability. This high probability extraction is crucial for their probabilistic argument to break the above uncertainty principle. Our main contribution to prove Theorem 1 is to show that, even with polynomial modulus, we can still argue the high success probability of their extractor. We first recall the distributions considered in [AHH24, Theorem 6]:

- $D^{\mathbf{A},\mathbf{y}}_{\mathsf{lwe}} = \left(\mathbf{A}, \mathbf{y}, \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal, \mathbf{s}^\intercal\mathbf{y} + e'\right)$ where $\mathbf{s} \leftarrow\!\!\$\ \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\chi}$, $e' \leftarrow \mathcal{D}_{\mathbb{Z},\chi'}$,
- $\bar{D}^{\mathbf{A},\mathbf{x}}_{\mathsf{lwe}} = \left(\mathbf{A}, \mathbf{Ax}, \mathbf{c}_0^\intercal, \mathbf{c}_0^\intercal\mathbf{x} + e''\right)$ where $\mathbf{c}_0^\intercal = \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal$, $e'' \leftarrow \mathcal{D}_{\mathbb{Z},\sqrt{\Sigma}}$,
- $D^{\mathbf{A},\mathbf{x}}_{\mathsf{gl}} = \left(\mathbf{A}, \mathbf{Ax}, \mathbf{c}_0^\intercal, \mathbf{c}_0^\intercal\mathbf{x} + e''\right)$ where $\mathbf{c}_0 \leftarrow\!\!\$\ \mathbb{Z}_q^m$,
- $D^{\mathbf{A},\mathbf{y}}_{\mathsf{unif}} = \left(\mathbf{A}, \mathbf{y}, \mathbf{c}_0^\intercal, c_1\right)$ where $\mathbf{c}_1 \leftarrow\!\!\$\ \mathbb{Z}_q$.

The goal is to argue the indistinguishability between $(D^{\mathbf{A},\mathbf{y}}_{\mathsf{lwe}}, \bar{D}^{\mathbf{A},\mathbf{x}}_{\mathsf{lwe}})$ and between $(\bar{D}^{\mathbf{A},\mathbf{x}}_{\mathsf{lwe}}, D^{\mathbf{A},\mathbf{x}}_{\mathsf{gl}})$, thereby obtaining a good distinguisher on $(D^{\mathbf{A},\mathbf{x}}_{\mathsf{gl}}, D^{\mathbf{A},\mathbf{y}}_{\mathsf{unif}})$, then use

this distinguisher to guess each element of $\mathbf{x}$ one by one via a search-to-decision reduction. To obtain the first step, they set $\Sigma = \chi'^2$ with $\chi'/\chi = \sigma\sqrt{m} \cdot 2^{o(n)}$ and invoke the noise smudging lemma, which results in the need for subexponential modulus. We achieve the same statistical indistinguishability with polynomial modulus by instead setting $\Sigma = \chi'^2 - \chi^2\|\mathbf{x}\|^2$ with $\chi'/\chi > \sigma\sqrt{2m}$, and we prove the following technical lemma.

Lemma 2. *$D^{\mathbf{A},\mathbf{y}}_{\mathsf{lwe}}$ and $\bar{D}^{\mathbf{A},\mathbf{x}}_{\mathsf{lwe}}$ have statistical distance at most negligible for any $\mathbf{x} \in \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2})$ and $\Sigma = \chi'^2 - \chi^2\|\mathbf{x}\|^2$ with $\chi'/\chi > \sigma\sqrt{2m}$.*

Proof. Given $(\mathbf{A}, \mathbf{x})$, such that $\mathbf{y} = \mathbf{A}\mathbf{x}$, we define intermediate distributions $(D_0, \ldots, D_4)$. These distributions have identically distributed first three components and differ only in their fourth component c_1. Below, $D^{\mathbf{A},\mathbf{A}\mathbf{x}}_{i,c_1}$ denotes the fourth component of D_i:

- $D^{\mathbf{A},\mathbf{A}\mathbf{x}}_{0,c_1} = \mathbf{s}^\intercal\mathbf{A}\mathbf{x} + e'$, where $e' \leftarrow \mathcal{D}_{\mathbb{Z},\chi'}$,
- $D^{\mathbf{A},\mathbf{A}\mathbf{x}}_{1,c_1} = \mathbf{s}^\intercal\mathbf{A}\mathbf{x} + \tilde{e} + \bar{e}$, where $\tilde{e} \leftarrow \mathcal{N}_{1,\chi'/\sqrt{2}}$ and $\bar{e} \leftarrow \mathcal{D}_{\mathbb{Z}-\tilde{e},\chi'/\sqrt{2}}$,
- $D^{\mathbf{A},\mathbf{A}\mathbf{x}}_{2,c_1} = \mathbf{s}^\intercal\mathbf{A}\mathbf{x} + \tilde{\mathbf{e}}_0\mathbf{x} + \tilde{e}_1 + \bar{e}$, where $\tilde{\mathbf{e}}_0 \leftarrow \mathcal{N}_{m,\chi\sqrt{2}}$, $\tilde{e}_1 \leftarrow \mathcal{N}_{1,\sqrt{\Sigma'}}$ with $\Sigma' = \chi'^2/2 - 2\chi^2\|\mathbf{x}\|^2$,
- $D^{\mathbf{A},\mathbf{A}\mathbf{x}}_{3,c_1} = \mathbf{s}^\intercal\mathbf{A}\mathbf{x} + (\mathbf{e} + \tilde{\mathbf{e}}'_0)\mathbf{x} + \tilde{e}_1 + \bar{e}$, where $\mathbf{e} \leftarrow \mathcal{N}_{m,\chi}, \tilde{\mathbf{e}}'_0 \leftarrow \mathcal{N}_{m,\chi}$,
- $D^{\mathbf{A},\mathbf{A}\mathbf{x}}_{4,c_1} = (\mathbf{s}^\intercal\mathbf{A} + \mathbf{e})\mathbf{x} + e''$, where $e'' \leftarrow \mathcal{D}_{\mathbb{Z},\sqrt{\Sigma}}$ with $\Sigma = \chi'^2 - \chi^2\|\mathbf{x}\|^2$.

By definition, we have $D_0 = D^{\mathbf{A},\mathbf{y}}_{\mathsf{lwe}}$ and $D_4 = \bar{D}^{\mathbf{A},\mathbf{x}}_{\mathsf{lwe}}$. We show that the statistical distance between D_0 and D_4 is negligible by proving that the statistical distance between D_{i,c_1} and D_{i+1,c_1} is negligible for $\forall\, i \in \{0,1,2,3\}$.

$D_{0,c_1} \approx_s D_{1,c_1}$ follows from Lemma 19 with any $\chi' \geq \sqrt{2}\eta_\varepsilon(\mathbb{Z})$ and $\varepsilon = 2^{-O(\lambda)}$, then the statistical distance between the two distributions is at most $4\varepsilon = \mathsf{negl}(\lambda)$.

$D_{1,c_1} \approx_s D_{2,c_1}$ follows from Lemma 21 with any $\chi > 0$ and $\chi' > 2 \cdot \chi \cdot \|\mathbf{x}\|$, then the two distributions are identical.

$D_{2,c_1} \approx_s D_{3,c_1}$ follows from Lemma 19 with any $\chi \geq \sqrt{2}\eta_\varepsilon(\mathbb{Z})$, and $\varepsilon = 2^{-O(\lambda)}$, then the statistical distance between the two distributions is at most $4\varepsilon = \mathsf{negl}(\lambda)$.

$D_{3,c_1} \approx_s D_{4,c_1}$ follows from Lemma 20 with the maximal eigenvalue of $\sqrt{\Sigma}, \lambda^{max}_{\sqrt{\Sigma}} \geq \eta_\varepsilon(\mathbb{Z}^m)$ and $\varepsilon = 2^{-O(\lambda)}$, then the statistical distance between e'' and $\tilde{\mathbf{e}}'_0\mathbf{x} + \tilde{e}_1 + \bar{e}$, hence between the two distributions, is at most $8\varepsilon = \mathsf{negl}(\lambda)$.

Combining these bounds, we have $D_{0,c_1} \approx_s D_{4,c_1}$ with statistical distance at most $16\varepsilon = \mathsf{negl}(\lambda)$, which concludes the proof. □

As Theorem 2 and Lemma 2 remove remaining subexponential modulus dependencies in [AHH24], we obtain the same result under polynomial modulus.

4 Identity-Based Encryption with SKL from LWE

In this section, we present a simple and efficient construction of selectively secure SKL-IBE with classical revocation from standard LWE, without relying on heavy machinery such as garbled circuits or obfuscation-based assumptions.

For our construction, we use the identity encoding scheme with full-rank differences (FRD) from [ABB10].

Definition 3 ([ABB10, Definition 20]). *Let q be a prime and $n \in \mathbb{N}$. We say that $\mathsf{H}_1 : \mathbb{Z}_q^n \to \mathbb{Z}_q^{n\times n}$ is an encoding with full-rank differences (FRD) if:*

1. *For all distinct $u, v \in \mathbb{Z}_q^n$, the matrix $\mathsf{H}_1(u) - \mathsf{H}_1(v) \in \mathbb{Z}_q^{n\times n}$ is full rank; and*
2. *H_1 is computable in polynomial time ($O(n \log q)$).*

Let $\mathsf{H}_2 : \{0,1\}^* \to \mathbb{Z}_q^n$ be a collision-resistant hash function. Defining $H := H_1 \circ H_2 : 0,1^* \to \mathbb{Z}_q^{n\times n}$ extends the identity encoding to arbitrary-length identities.

4.1 Construction

Construction 2. Let λ be the security parameter. Let $q \geq 2$ be a prime, let $n, m, \nu \in \mathbb{N}$ and $\sigma, \chi, \chi', \beta_0 > 0$ be parameters. Our construction for a SKL-IBE with classical revocation scheme $\mathsf{SKL\text{-}IBE} = (\mathsf{Setup}, \mathcal{KG}, \mathsf{Enc}, \mathcal{Dec}, \mathcal{Del}, \mathsf{Verif})$ uses the following building blocks:

- A post-quantum secure pseudorandom function $\mathsf{F} : \{0,1\}^\lambda \times \mathbb{Z}_q^{2n} \to \{0,1\}^m$.
- An FRD identity encoding scheme $\mathsf{H} : \{0,1\}^* \to \mathbb{Z}_q^{n\times n}$ by combining H_1 and H_2 as in Definition 3.

The construction is defined as follows:

- $\mathsf{Setup}(1^\lambda) \to (\mathsf{mpk}, \mathsf{msk})$:
  - Sample $(\mathbf{A}, \mathbf{T}_\mathbf{A}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$, $\bar{\mathbf{A}}, \mathbf{B} \leftarrow\!\$\ \mathbb{Z}_q^{n\times m}$, and $\mathbf{y} \leftarrow\!\$\ \mathbb{Z}_q^n$.
  - Generate PRF key $\mathsf{k} \leftarrow\!\$\ \{0,1\}^\lambda$.
  - Output $\mathsf{mpk} = (\mathbf{A}, \bar{\mathbf{A}}, \mathbf{B}, \mathbf{y})$ and $\mathsf{msk} = (\mathbf{T}_\mathbf{A}, \mathsf{k})$.
- $\mathcal{KG}(\mathsf{msk}, \mathsf{id}) \to \rho_\mathsf{id}$:
  - For $\mathbf{A}_\mathsf{id} = [\mathbf{A} | \bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B}]$, compute $\mathbf{T}_{\mathbf{A}_\mathsf{id}} \leftarrow \mathsf{ExtendRight}(\mathbf{A}, \bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B}, \mathbf{T}_\mathbf{A})$.
  - Generate a Gaussian coset state $|\psi_{\mathbf{A}_\mathsf{id},\mathbf{y}}\rangle \leftarrow \mathit{SampGauss}(\mathbf{A}_\mathsf{id}, \mathbf{T}_{\mathbf{A}_\mathsf{id}}, \mathbf{y}, \sigma)$

$$\rho = |\psi_{\mathbf{A}_\mathsf{id},\mathbf{y}}\rangle = \sum_{\bar{\mathbf{x}} \in \mathbb{Z}_q^{2m} : \mathbf{A}_\mathsf{id}\bar{\mathbf{x}} = \mathbf{y}} \rho_\sigma(\bar{\mathbf{x}}) \, |\bar{\mathbf{x}}\rangle \, .$$

  - Measure the last m coordinates of ρ in the computational basis to obtain

$$\rho' = \left|\psi_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle \otimes |\mathbf{u}\rangle = \sum_{\mathbf{x} \in \mathbb{Z}_q^m : \mathbf{A}\mathbf{x} = \hat{\mathbf{y}}} \rho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \otimes |\mathbf{u}\rangle \, ,$$

    where $\mathbf{u} \in \mathbb{Z}_q^m$ is the measurement outcome and some $\hat{\mathbf{y}} \in \mathbb{Z}_q^n$.
  - Compute $\mathbf{v} \leftarrow \mathsf{F}(\mathsf{k}, \mathsf{H}_2(\mathsf{id}) \| \hat{\mathbf{y}})$ and the state $\left|\psi^\mathbf{v}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle = \mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}} \, |\psi_{\mathbf{A},\hat{\mathbf{y}}}\rangle \, .$
  - Output $\rho_\mathsf{id} = \left(\left|\psi^\mathbf{v}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle, \hat{\mathbf{y}}, \mathbf{u} \right)$.
- $\mathsf{Enc}(\mathsf{mpk}, \mathsf{id} \in \{0,1\}^*, \mu \in \{0,1\}) \to \mathsf{ct}_\mu$: Sample a secret $\mathbf{s} \leftarrow\!\$\ \mathbb{Z}_q^n$ and errors $\mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^m,\chi}, e' \leftarrow \mathcal{D}_{\mathbb{Z},\chi'}$. Sample a random $\mathbf{R} \leftarrow\!\$\ \{-1,1\}^{m\times m}$. For $\mathbf{A}_\mathsf{id} = [\mathbf{A} \mid \bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B}]$, output:

$$\mathsf{ct}_{\mathsf{id},\mu} = \left(\mathbf{s}^\intercal \mathbf{A}_\mathsf{id} + \left[\mathbf{e}_0^\intercal \mid \mathbf{e}_0^\intercal \mathbf{R}\right], \mathbf{s}^\intercal \mathbf{y} + e' + \mu \cdot \lfloor \tfrac{q}{2} \rfloor\right) \in \mathbb{Z}_q^{2m+1}.$$

– $\mathcal{Dec}(\rho_{\mathsf{id}}, \mathsf{ct}) \to \{0,1\}$: Apply the unitary $U : |\mathbf{x}\rangle|\mathbf{0}\rangle \to |\mathbf{x}\rangle|\mathsf{ct}\cdot(-\mathbf{x},-\mathbf{u},1)\rangle$ on input $\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle|\mathbf{0}\rangle$. Measure the last register in the computational basis. Output 0 if the outcome is closer to 0 than to $\lfloor q/2 \rfloor$, and 1 otherwise.
– $\mathcal{Del}(\rho_{\mathsf{id}}) \to \mathsf{cert}$: Apply the m-qudit q-ary quantum Fourier transform to the state ρ_{id}, then measure in the computational basis to obtain an outcome $\mathbf{w} \in \mathbb{Z}_q^m$. Output $\mathsf{cert} := (\mathbf{w}, \hat{\mathbf{y}}, \mathsf{id})$.
– $\mathsf{Verif}(\mathsf{msk}, \mathsf{cert}) \to \{\top, \bot\}$:
  • Parse $\mathsf{msk} = (\mathbf{T_A} \in \mathbb{Z}_q^{m\times m}, \mathsf{k})$, $\mathsf{cert} = (\mathbf{w}, \hat{\mathbf{y}}, \mathsf{id})$.
  • Compute $\mathbf{v} \leftarrow \mathsf{F}(\mathsf{k}, \mathsf{H}_2(\mathsf{id}) \| \hat{\mathbf{y}})$.
  • Compute $\mathbf{c}^\intercal = \mathbf{w}^\intercal \cdot \mathbf{T_A} \pmod q$ and $\mathbf{d}^\intercal = \mathbf{c}^\intercal \cdot \mathbf{T_A}^{-1}$, where $\mathbf{T_A}^{-1} \in \mathbb{R}^{m\times m}$.
  • For $i \in [m]$, let $v'_i = 0$, if $d_i \in [-q/\sigma, q/\sigma]$, and let $v'_i = 1$, otherwise.
  • Output $\top$, if $\mathbf{v}' = \mathbf{v}$, and output $\bot$ otherwise.

Parameters. Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n, m \in \mathbb{N}$, q be a prime modulus with $m = \lceil 6n \log q \rceil$ and $\sqrt{8m} < \sigma < q/\sqrt{8m}$, and $\nu = 64m^2$, each parameterized by λ. Let $\chi, \chi' > 2\sqrt{n}$ such that $\chi'/\chi > \mathcal{O}(\sigma m)$. Let $\beta_0 = \mathcal{O}(\chi\sigma m^{3/2} + \chi')\lambda^{\omega(1)}$ be the noise bound for decryption correctness.

Correctness. For $\mathsf{id} \in \{0,1\}^*$ and $\bar{\mathbf{x}}^\intercal := [\mathbf{x}^\intercal \mid \mathbf{u}^\intercal] \in \mathbb{Z}_q^{2m}$, we have $\mathbf{A}_{\mathsf{id}}\bar{\mathbf{x}} = \mathbf{A}\mathbf{x} + \big(\bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B}\big)\mathbf{u}$. Therefore, for $(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda)$, $\left(\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle, \hat{\mathbf{y}}, \mathbf{u}\right) \leftarrow \mathcal{KG}(\mathsf{msk}, \mathsf{id})$, we have $\hat{\mathbf{y}} + \big(\bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B}\big)\mathbf{u} = \mathbf{y}$. For $\mu \in \{0,1\}$ and $\mathsf{id} \in \{0,1\}^*$, we have:

$$\begin{aligned}
|\mathsf{ct}\cdot(-\mathbf{x},-\mathbf{u},1)\rangle &= \left|\left(\mathbf{s}^\intercal\left[\mathbf{A} \mid \bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B}\right] + \left[\mathbf{e}_0^\intercal \mid \mathbf{e}_0^\intercal\mathbf{R}\right], \mathbf{s}^\intercal\mathbf{y} + e' + \mu\cdot\lfloor\tfrac{q}{2}\rfloor\right)\cdot(-\mathbf{x},-\mathbf{u},1)\right\rangle \\
&= \left|-\mathbf{s}^\intercal\mathbf{A}\mathbf{x} - \mathbf{e}_0^\intercal\mathbf{x} - \mathbf{s}^\intercal\big(\bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B}\big)\mathbf{u} - \mathbf{e}_0^\intercal\mathbf{R}\mathbf{u} + \mathbf{s}^\intercal\mathbf{y} + e' + \mu\cdot\lfloor\tfrac{q}{2}\rfloor\right\rangle \\
&= \left|\mathbf{s}^\intercal(-\mathbf{A}\mathbf{x} - (\bar{\mathbf{A}} + \mathsf{H}(\mathsf{id})\mathbf{B})\mathbf{u} + \mathbf{y}) - \mathbf{e}_0^\intercal\mathbf{x} - \mathbf{e}_0^\intercal\mathbf{R}\mathbf{u} + e' + \mu\cdot\lfloor\tfrac{q}{2}\rfloor\right\rangle \\
&= \left|-\mathbf{e}_0^\intercal\mathbf{x} - \mathbf{e}_0^\intercal\mathbf{R}\mathbf{u} + e' + \mu\cdot\lfloor\tfrac{q}{2}\rfloor\right\rangle \\
&\propto |\ \mu \pmod q\rangle\,.
\end{aligned}$$

With overwhelming probability, we have $\|\mathbf{x}\| \le \sigma\sqrt{m}$ and $\|\mathbf{u}\| \le \sigma\sqrt{m}$, thus:

$$\| - \mathbf{e}_0^\intercal\mathbf{x} - \mathbf{e}_0^\intercal\mathbf{R}\mathbf{u} + e'\| \le \mathcal{O}(\chi\sigma m^{3/2} + \chi') < \beta_0.$$

As long as β_0 is bounded by $q/4$, decryption correctness follows with overwhelming probability. Reusability then follows immediately from the gentle measurement lemma [Win99], while verification correctness follows immediately from [APV23].

4.2 Security Proof

Theorem 3. *Assuming the quantum hardness of* LWE *with polynomial modulus, the scheme* SKL-IBE $= (\mathsf{Setup}, \mathcal{KG}, \mathsf{Enc}, \mathcal{Dec}, \mathcal{Del}, \mathsf{Verif})$ *in Construction 2 with specified parameters is secure according to Definition 2.*

Proof. We first assume the existence of an adversary $\mathcal{A}$ that breaks the IND-SKL security (Figure 1) of our SKL-IBE construction given public parameters and a master public key $(\mathbf{A}, \bar{\mathbf{A}}, \mathbf{B}, \mathbf{y})$. Then, we construct two reduction adversaries

$(\mathcal{R}_0, \mathcal{R}_1)$ with non-negligible advantage in the uncertainty principle experiment in Theorem 2 w.r.t challenge $\left(\left|\psi^{\mathbf{v}_{\mathsf{id}^*}}_{\mathbf{A},\hat{\mathbf{y}}_{\mathsf{id}^*}}\right\rangle, \hat{\mathbf{y}}_{\mathsf{id}^*}\right)$, where id^* is the challenge identity in the IND-SKL game. We first consider the following sequence of security games:

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

1. $\mathcal{A}$ sends a challenge identity id^* to $\mathcal{CH}$.
2. $\mathcal{CH}$ samples $(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda)$. Then, $\mathcal{CH}$ sends mpk to $\mathcal{A}$.
3. $\mathcal{A}$ sends queries to the key generation oracle $\mathcal{O}_{\mathcal{KG}}$ to receive $\rho_{\mathsf{id}_i} := \left(\left|\psi^{\mathbf{v}_{\mathsf{id}_i}}_{\mathbf{A},\hat{\mathbf{y}}_{\mathsf{id}_i}}\right\rangle, \hat{\mathbf{y}}_{\mathsf{id}_i}, \mathbf{u}_{\mathsf{id}_i}\right)$ with at most one query with $\mathsf{id}_i = \mathsf{id}^*$.
4. $\mathcal{CH}$ sends REVOKE to $\mathcal{A}$, then $\mathcal{A}$ returns classical certificates $\mathsf{cert}_{\mathsf{id}^*} := (\mathbf{w}_{\mathsf{id}^*}, \hat{\mathbf{y}}_{\mathsf{id}^*}, \mathsf{id}^*)$.
5. $\mathcal{CH}$ continues if $\mathsf{Verif}(\mathsf{msk}, \mathsf{cert}_{\mathsf{id}^*}) \to \top$, otherwise returns $\perp$ and aborts.
6. $\mathcal{A}$ sends the plaintext bit μ to $\mathcal{CH}$.
7. If $b = 0$, $\mathcal{CH}$ sends $\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{mpk}, \mathsf{id}^*, \mu)$ to $\mathcal{A}$. Otherwise, if $b = 1$, $\mathcal{CH}$ sends $\mathsf{ct} \leftarrow\!\!\$\ \mathcal{C}$, where $\mathcal{C}$ is the ciphertext space.
8. $\mathcal{A}$ outputs a guess b', $\mathcal{CH}$ outputs b' as the output of the experiment.

Fig. 1: Security experiment $\mathsf{Expt}^{\mathsf{SKL\text{-}IBE}}_{\mathcal{A}}(1^\lambda, b)$.

$\mathsf{G}_0(\lambda)$**:** This is the original game $\mathsf{Expt}^{\mathsf{SKL\text{-}IBE}}_{\mathcal{A}}(1^\lambda, b)$ as defined in Figure 1.
$\mathsf{G}_1(\lambda)$**:** We replace the pseudorandom function F with a random function G.
$\mathsf{G}_2(\lambda)$**:** We sample $\mathbf{R}^*$ used for encryption at setup and set $\bar{\mathbf{A}} := \mathbf{A}\mathbf{R}^* - \mathsf{H}(\mathsf{id}^*)\mathbf{B}$.
$\mathsf{G}_3(\lambda)$**:** We sample $\left(\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle, \hat{\mathbf{y}}\right) \leftarrow \mathcal{GenGauss}(\mathbf{A}, \sigma)$ for $\mathbf{v} \leftarrow\!\!\$\ \{0,1\}^m$, $\mathbf{u} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \sigma/\sqrt{2}}$, and set $\mathbf{y} := \hat{\mathbf{y}} + \left(\bar{\mathbf{A}} + \mathsf{H}(\mathsf{id}^*)\mathbf{B}\right)\mathbf{u}$ when generating mpk. For the key query with id^*, the key generation oracle returns $\rho_{\mathsf{id}^*} := \left(\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle, \hat{\mathbf{y}}, \mathbf{u}\right)$.
$\mathsf{G}_4(\lambda)$**:** We sample $\mathbf{A} \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times m}$ and $(\mathbf{B}, \mathbf{T}_{\mathbf{B}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$ at setup.
- For the verification procedure, replace the efficient algorithm $\mathsf{Verif}(\cdot)$ with an inefficient algorithm $\mathsf{IneffVrfy}(\cdot)$ that given $\mathbf{A}$, computes a trapdoor $\mathbf{T}'_{\mathbf{A}}$ in exponential time and runs $\mathsf{Verif}(\cdot)$ with $\mathbf{T}'_{\mathbf{A}}$.
- For the key generation procedure, we compute $\mathbf{T}_{\mathbf{A}_{\mathsf{id}}} \leftarrow \mathsf{ExtendLeft}(\mathbf{A}, (\mathsf{H}(\mathsf{id}) - \mathsf{H}(\mathsf{id}^*))\,\mathbf{B}, \mathbf{T}_{\mathbf{B}}, \mathbf{R}^*)$ and the rest are the same as in $\mathcal{KG}$.

Lemma 3. G_0 *and* G_1 *are computationally indistinguishable.*

Proof. Indistinguishability follows directly from the pseudorandomness of F. □

Lemma 4. G_1 *and* G_2 *are statistically indistinguishable.*

Proof. This follows from Lemma 18, as the distribution $(\mathbf{A}, \mathbf{A}\mathbf{R}^*, \mathbf{e}_0^\intercal \mathbf{R}^*)$ is statistically close to the distribution $(\mathbf{A}, \mathbf{A}', \mathbf{e}_0^\intercal \mathbf{R}^*)$, where $\mathbf{A}' \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times m}$. □

Lemma 5. G_2 *and* G_3 *are statistically indistinguishable.*

Proof. This follows from the definition of the Gaussian coset state $\left|\psi_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle$, where the distribution of $\hat{\mathbf{y}}$ is statistically close to the uniform distribution over $\mathbb{Z}_q^n$. □

Lemma 6. G_3 *and* G_4 *are statistically indistinguishable.*

Proof. We equip the inefficient verification oracle with an inefficient algorithm $\mathsf{IneffVrfy}$ that is allowed to run in exponential time. Since the original trapdoor $\mathbf{T_A}$ is no longer necessary, the switch is statistically indistinguishable (Lemma 22). For key generation queries, since we can answer $\mathsf{id} = \mathsf{id}^*$ without any trapdoor as in G_3, thus we only need to analyze the case for $\mathsf{id} \neq \mathsf{id}^*$ where we have $\mathbf{A}_{\mathsf{id}} = [\mathbf{A} \mid \mathbf{A}\mathbf{R}^* + (\mathsf{H}(\mathsf{id}) - \mathsf{H}(\mathsf{id}^*))\mathbf{B}]$. By definition of the FRD encoding scheme (Definition 3), the matrix $\mathsf{H}(\mathsf{id}) - \mathsf{H}(\mathsf{id}^*)$ is full rank. Then by Lemma 22 and any $\mathbf{R} \in \{-1, 1\}^{m \times m}$, the trapdoor $\mathbf{T}_{\mathbf{A}_{\mathsf{id}}} \leftarrow \mathsf{ExtendLeft}(\mathbf{A}, (\mathsf{H}(\mathsf{id}) - \mathsf{H}(\mathsf{id}^*))\mathbf{B}, \mathbf{T_B}, \mathbf{R})$ also allows us to invert the LWE problem for the matrix $\mathbf{A}_{\mathsf{id}}$, which can be used to answer key generation queries for $\mathsf{id} \neq \mathsf{id}^*$ as in G_3. □

Lemma 7. *Assuming the quantum hardness of* LWE*, we have* $\mathsf{Adv}_4(\mathcal{A}, \mathcal{D}) \leq \mathsf{negl}(\lambda)$.

Proof. We assume that there exists a QPT adversary $\mathcal{A}$ that breaks the security game G_4 with some non-negligible, in particular inverse polynomial, advantage

$$\mathsf{Adv}_4(\mathcal{A}, \mathcal{D}) = \left|\Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{G}_4}_{\mathcal{A},\mathcal{D}}(1^\lambda, 0)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{G}_4}_{\mathcal{A},\mathcal{D}}(1^\lambda, 1)\right]\right| = \varepsilon(\lambda) \geq \frac{1}{\mathsf{poly}(\lambda)}.$$

$\underline{\mathsf{Expt}^{\mathsf{G}_4}_{\mathcal{A},\mathcal{D}}(1^\lambda, b)}$:

1. Generate $\mathsf{id}^* \leftarrow \mathcal{A}(1^\lambda)$.
2. Sample $(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda)$.
3. Generate $\rho \otimes \tau_\lambda \leftarrow \mathcal{A}^{\mathcal{O}_{\mathcal{KG}}}(\mathsf{mpk})$, τ_λ is any polynomial-sized advice state.
4. Generate $|\mathsf{cert}_{\mathsf{id}^*}\rangle\langle\mathsf{cert}_{\mathsf{id}^*}| \otimes \rho_{\mathsf{aux}} \leftarrow \mathcal{A}(\mathsf{mpk}, \rho \otimes \tau_\lambda)$.
5. Continue if $\mathsf{IneffVrfy}(\mathsf{msk}, \mathsf{cert}_{\mathsf{id}^*}) \to \top$, otherwise abort with output $\bot$.
6. Sample $\mathsf{ct} \leftarrow D_b$, where $D_0 \sim (\mathbf{c} \leftarrow \mathsf{Enc}(\mathsf{mpk}, \mathsf{id}^*, 0))$ and $D_1 \sim (\mathbf{c} \leftarrow\!\!\$\, \mathcal{C})$.
7. Generate $b' \leftarrow \mathcal{D}^{\mathcal{O}_{\mathcal{KG}}}(\mathsf{mpk}, \mathsf{ct}_b, \rho_{\mathsf{aux}})$.
8. Output b' as the output of the experiment.

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

1. Generate $\mathsf{id}^* \leftarrow \mathcal{A}(1^\lambda)$.
2. Sample $(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda)$.
3. Generate $\rho \otimes \tau_\lambda \leftarrow \mathcal{A}^{\mathcal{O}_{\mathcal{KG}}}(\mathsf{mpk})$, τ_λ is any polynomial-sized advice state.
4. Generate $|\mathsf{cert}_{\mathsf{id}^*}\rangle\langle\mathsf{cert}_{\mathsf{id}^*}| \otimes \rho_{\mathsf{aux}} \leftarrow \mathcal{A}(\mathsf{mpk}, \rho \otimes \tau_\lambda)$.
5. Continue if $\mathsf{IneffVrfy}(\mathsf{msk}, \mathsf{cert}_{\mathsf{id}^*}) \to \top$, otherwise abort with output $\bot$.
6. Output $b \leftarrow \mathsf{ATI}^{\frac{\gamma}{6},\delta}_{\mathcal{P}_{\mathcal{D}},(D_0,D_1),\frac{1}{2}+\gamma}(\rho_{\mathsf{aux}})$ where $\delta = 2^{-\Theta(\lambda)}$.

Fig. 2: Implement distinguisher in $\mathsf{Expt}^{\mathsf{G}_4}$ with ATI.

We assume that the adversary submits $\mu = 0$ w.l.o.g. In Figure 2, we simplify the security experiment of G_4 and define a separated distinguisher $\mathcal{D}$ that takes as input the residual state ρ_{aux} after revocation. Given a quantum extractor $\mathcal{E}$ that, on input $(\mathcal{D}, \rho_{\mathsf{aux}})$, outputs a vector $\mathbf{x} \in \mathbb{Z}_q^m$, we define the following events:

– $\mathsf{VrfySuc}$: the verification algorithm $\mathsf{IneffVrfy}$ outputs $\top$.
– $\mathsf{GoodDec}$: the distinguisher $\mathcal{D}$ outputs $b' = b$.
– ExtSuc: the extractor $\mathcal{E}$ outputs a short $\mathbf{x} : \mathbf{A}\mathbf{x} = \hat{\mathbf{y}}_{\mathsf{id}^*}$, with $\hat{\mathbf{y}}_{\mathsf{id}^*}$ from $\mathsf{cert}_{\mathsf{id}^*}$.

If $\mathcal{A}$ wins G_4 with advantage $\varepsilon(\lambda)$, then $\Pr[\mathsf{VrfySuc} \wedge \mathsf{GoodDec}] \geq 1/\mathsf{poly}(\lambda)$. The proof proceeds in two main steps:

1. We construct an extractor $\mathcal{E}$ satisfying $\Pr[\mathsf{ExtSuc} \mid \mathsf{GoodDec}] \geq 1 - \mathsf{negl}(\lambda)$. From probabilistic Claim B.1, we have $\Pr[\mathsf{VrfySuc} \wedge \mathsf{ExtSuc}] \geq 1/\mathsf{poly}(\lambda)$.
2. From such adversary and extractor, we construct a reduction to break the uncertainty principle of Gaussian coset states (Theorem 2).

To estimate the success probability of the quantum distinguisher $\mathcal{D}$, we apply an Approximate Threshold Implementation (ATI) (Lemma 23) to efficiently test the *goodness* of $\mathcal{D}$ while only negligibly disturbing the residual state. For distributions D_0 and D_1, let (D_0, D_1) denote the distribution of (b, ct) obtained by sampling

$$b \leftarrow\!\$\ \{0,1\} \qquad \text{and} \qquad \mathsf{ct} \leftarrow D_b.$$

The corresponding ATI procedure tests whether the distinguisher, given its auxiliary state, can predict b from ct with probability at least $1/2 + \gamma$.

Lemma 8. *For adversary $(\mathcal{A}, \mathcal{D})$ such that $\mathsf{Adv}_4(\mathcal{A}, \mathcal{D}) = \mu(\lambda)$, it satisfies*

$$\Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{ATI}}_{\mathcal{A},\mathcal{D},\gamma}(1^\lambda)\right] \geq \frac{\mu(\lambda)}{4} - \mathsf{negl}(\lambda)$$

for $\gamma = \frac{3\mu}{14}$ where $\mathsf{Expt}^{\mathsf{ATI}}_{\mathcal{A},\mathcal{D},\gamma}(1^\lambda)$ is shown in Figure 2.

Proof. The proof of Lemma 8 is identical to the proof of [AHH24, Lemma 13], so we refer readers to the original proof for details. □

For coherence, we postpone the analysis of the extractor to Section 4.3 and only state the following result.

Lemma 9. *Assuming the quantum hardness of LWE with suitable parameters, there exists a QPT extractor $\mathcal{E}$ that takes as input the distinguisher $\mathcal{D}$ and the residual state ρ'_{aux} satisfies*

$$\Pr\left[\begin{array}{c} \mathbf{x} \in \Lambda_q^{\hat{\mathbf{y}}_{\mathsf{id}^*}}(\mathbf{A}) \\ \cap \\ \mathcal{B}^m(0, \frac{\sigma\sqrt{m}}{2}) \end{array} : \begin{array}{c} \mathsf{id}^* \leftarrow \mathcal{A}(1^\lambda), \\ (\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda), \\ \rho \otimes \tau_\lambda \leftarrow \mathcal{A}^{\mathcal{O}_{\mathcal{KG}}}(\mathsf{mpk}), \\ |\mathsf{cert}_{\mathsf{id}^*}\rangle\langle\mathsf{cert}_{\mathsf{id}^*}| \otimes \rho_{\mathsf{aux}} \leftarrow \mathcal{A}^{\mathcal{O}_{\mathcal{KG}}}(\mathsf{mpk}, \rho \otimes \tau_\lambda), \\ 1 \leftarrow \mathsf{ATI}^{\frac{\gamma}{6},\delta}_{\mathcal{P}_{\mathcal{D}},(D_0,D_1),\frac{1}{2}+\gamma}(\rho_{\mathsf{aux}}), \\ \mathbf{x} \leftarrow \mathcal{E}(\mathcal{D}, \rho'_{\mathsf{aux}}) \end{array}\right] \geq 1 - \mathsf{negl}(\lambda),$$

where $\delta = 2^{-\Theta(\lambda)}$.

We proceed to show how to construct our reduction adversaries $(\mathcal{R}_0, \mathcal{R}_1)$ to win the uncertainty principle experiment $\mathsf{UncertaintyExpt}(1^\lambda)$ from Theorem 2.

– $\mathcal{R}_0(\mathbf{A}, \mathbf{y}, |\psi^{\mathbf{v}}_{\mathbf{A},\mathbf{y}}\rangle)$:
  1. Run the algorithm $\mathcal{A}$ to generate a challenge identity id^*.
  2. Sample $\bar{\mathbf{A}}, \mathbf{B}, \mathbf{R}$ and simulate the key generation oracle $\mathcal{O}_{\mathcal{KG}}$ as in G_4, then run the algorithm $\mathcal{A}$ to generate $\rho \otimes \tau_\lambda \leftarrow \mathcal{A}^{\mathcal{O}_{\mathcal{KG}}}(\mathsf{mpk})$.
  3. Continue running $\mathcal{A}$ to get $|\mathsf{cert}_{\mathsf{id}^*}\rangle\langle\mathsf{cert}_{\mathsf{id}^*}| \otimes \rho_{\mathsf{aux}} \leftarrow \mathcal{A}^{\mathcal{O}_{\mathcal{KG}}}(\mathsf{mpk}, \rho \otimes \tau_\lambda)$.
  4. Run $\mathsf{ATI}^{\frac{\gamma}{6},\delta}_{\mathcal{P}_{\mathcal{D}},(D_0,D_1),\frac{1}{2}+\gamma}(\rho_{\mathsf{aux}})$ and abort if the output is 0.
  5. Continue to run the extractor $\mathcal{E}$ to obtain $\mathbf{x} \leftarrow \mathcal{E}(\mathcal{D}, \rho'_{\mathsf{aux}})$.
  6. Send $\mathbf{x}$ to $\mathcal{Ch}$ and initialize $\mathcal{R}_1$ with $(\mathbf{A}, \mathbf{w}_{\mathsf{id}^*})$.
– $\mathcal{R}_1(\mathbf{A}, \mathbf{w}_{\mathsf{id}^*})$: Run the inefficient verification algorithm $\mathsf{IneffVrfy}$ to compute a trapdoor $\mathbf{T}'_{\mathbf{A}}$, then use it to recover $\mathbf{v}'$ and send to $\mathcal{Ch}$.

With the results of Lemmas 8 and 9 and Claim B.1, we argue that the $(\mathcal{R}_0, \mathcal{R}_1)$ described above win the uncertainty principle experiment $\mathsf{UncertaintyExpt}(1^\lambda)$ with inverse polynomial probability, thus we complete the proof. □

Combining Lemmas 3 to 7, we conclude the proof of Theorem 3. □

4.3 Quantum Extractor

In this section, we revisit the black-box quantum extractor from Lemma 9 and show that $\Pr[\mathsf{ExtSuc} \mid \mathsf{GoodDec}] \geq 1 - \mathsf{negl}(\lambda)$. Thus we may drop the inefficient verification procedure $\mathsf{IneffVrfy}$ in Figure 2 and assume the adversary is a good quantum distinguisher for the challenge distributions, i.e., ATI outputs 1. Then, we construct a quantum extractor that succeeds with probability $1 - \mathsf{negl}(\lambda)$ by adapting the quantum *search-to-decision* extractor of [AHH24] with additional key query oracle for the SKL-IBE adversary. Moreover, our distinguisher operates on different distributions:

– $D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}} = \left(\mathbf{A}, \mathbf{AR}, \hat{\mathbf{y}}, \mathbf{u}, \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}_0^\intercal, (\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}_0^\intercal)\mathbf{R}, \mathbf{s}^\intercal\mathbf{y} + e'\right)$ where $\mathbf{s} \leftarrow\!\$\, \mathbb{Z}_q^n$, $\mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^m,\chi}$, $e' \leftarrow \mathcal{D}_{\mathbb{Z},\chi'}$ with $\mathbf{y} = \mathbf{ARu} + \hat{\mathbf{y}}$,
– $\bar{D}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{lwe}} = \left(\mathbf{A}, \mathbf{AR}, \mathbf{Ax}, \mathbf{u}, \mathbf{c}_0^\intercal, \mathbf{c}_0^\intercal\mathbf{R}, \mathbf{c}_0^\intercal(\mathbf{x} + \mathbf{Ru}) + e''\right)$ where $\mathbf{c}_0^\intercal = \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}_0^\intercal$, $e'' \leftarrow \mathcal{D}_{\mathbb{Z},\sqrt{\Sigma''}}$ with $\Sigma'' = \chi'^2 - \chi^2\|\mathbf{x} + \mathbf{Ru}\|^2$,
– $D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}} = \left(\mathbf{A}, \mathbf{AR}, \mathbf{Ax}, \mathbf{u}, \mathbf{c}_0^\intercal, \mathbf{c}_0^\intercal\mathbf{R}, \mathbf{c}_0^\intercal(\mathbf{x} + \mathbf{Ru}) + e''\right)$ where $\mathbf{c}_0 \leftarrow\!\$\, \mathbb{Z}_q^m$,
– $D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}} = \left(\mathbf{A}, \mathbf{AR}, \mathbf{Ax}, \mathbf{u}, \mathbf{c}_0^\intercal, \mathbf{c}_0^\intercal\mathbf{R}, c_2\right)$ where $\mathbf{c}_0 \leftarrow\!\$\, \mathbb{Z}_q^m, c_2 \leftarrow\!\$\, \mathbb{Z}_q$,
– $D^{\mathbf{A},\mathbf{A}',\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{unif}} = \left(\mathbf{A}, \mathbf{A}', \hat{\mathbf{y}}, \mathbf{u}, \mathbf{c}_0^\intercal, \mathbf{c}_1^\intercal, c_2\right)$ where $\mathbf{c}_1 \leftarrow\!\$\, \mathbb{Z}_q^m$ with $\mathbf{y} = \mathbf{A}'\mathbf{u} + \hat{\mathbf{y}}$.

For each of these distributions, the modification $D(i, g_i)$ is defined as

$$D(i, g_i) = (\mathbf{A}_0, \mathbf{A}_1, \mathbf{v}_0, \mathbf{v}_1, \mathbf{c}_0 + c \cdot \hat{\mathbf{i}}, \mathbf{c}_1 + c \cdot \mathbf{r}_i, c_2 + c \cdot (g_i + (\mathbf{Ru})_i)),$$

where $c \leftarrow\!\$\, \mathbb{Z}_q$, $\hat{\mathbf{i}}$ is the standard unit vector, i.e., the entry at the index i-th is 1 and all other entries are 0, $\mathbf{r}_i$ is the i-th row of the matrix $\mathbf{R}$ and $(\mathbf{Ru})_i$ is the i-th entry of the vector $\mathbf{Ru}$. We proceed to prove some useful lemmas.

Lemma 10. *$D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}}$ and $\bar{D}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{lwe}}$ have statistical distance at most negligible for any $\mathbf{x}, \mathbf{u} \in \mathcal{B}^m(\mathbf{0}, \sigma\sqrt{m/2})$, $\mathbf{R} \in \{-1, 1\}^{m\times m}$, and $\chi'/\chi > \mathcal{O}(\sigma m)$.*

Proof. The proof for this lemma is the same as the proof of Lemma 2, except that instead of $\mathbf{A}, \mathbf{x}$ and $\mathbf{y} = \mathbf{A}\mathbf{x}$, we consider $\mathbf{A}, \mathbf{f}$ where $\mathbf{f} = \mathbf{x} + \mathbf{R}\mathbf{u}$ and $\mathbf{y} = \mathbf{A}\mathbf{f}$. □

Lemma 11. *The following holds for any integer $i \in [m]$:*

$$\begin{cases} D_{\mathsf{cg}}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}(i, g_i) = D_{\mathsf{cg}}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}} & \text{if } g_i = x_i, \quad (2) \\ D_{\mathsf{cg}}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}(i, g_i) = D_{\mathsf{wg}}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}} & \text{if } g_i \neq x_i. \quad (3) \end{cases}$$

Proof. By definition, $D(i, g_i)$ only modifies the three values $\mathbf{c}_0$, $\mathbf{c}_1$, and c_2, thus we omit the other values in the distributions for simplicity and we have:

$$\begin{aligned} D_{\mathsf{cg}}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}(i, g_i) &= (\mathbf{c}_0^\intercal + c \cdot \hat{\mathbf{i}}^\intercal, \mathbf{c}_0^\intercal \mathbf{R} + c \cdot \mathbf{r}_i, \mathbf{c}_0^\intercal(\mathbf{x} + \mathbf{R}\mathbf{u}) + e'' + c \cdot (g_i + (\mathbf{R}\mathbf{u})_i)) \\ &= (\mathbf{c}_0^\intercal + c \cdot \hat{\mathbf{i}}^\intercal, (\mathbf{c}_0^\intercal + c \cdot \hat{\mathbf{i}}^\intercal)\mathbf{R}, (\mathbf{c}_0^\intercal + c \cdot \hat{\mathbf{i}}^\intercal)(\mathbf{x} + \mathbf{R}\mathbf{u}) + e'' + c \cdot (g_i - \mathbf{x}_i)). \end{aligned}$$

We can see that the distribution is preserved if $g_i = x_i$, thus we have Equation (2). And if $g_i \neq x_i$, c_2 is uniform over $\mathbb{Z}_q$ since $c \leftarrow\!\!\$\, \mathbb{Z}_q$ and q is prime, thus we have Equation (3). □

Lemma 12. *$D_{\mathsf{wg}}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}$ and $D_{\mathsf{unif}}^{\mathbf{A},\mathbf{A}',\hat{\mathbf{y}},\mathbf{y}}$ have statistical distance at most negligible for any $\mathbf{R} \in \{-1, 1\}^{m \times m}$.*

Proof. This follows from the generalized leftover hash lemma (Lemma 18) for the uniform matrix $\begin{bmatrix} \mathbf{A} \\ \mathbf{c}_0^\intercal \end{bmatrix}$ so that we have the distribution of $(\mathbf{A}, \begin{bmatrix} \mathbf{A} \\ \mathbf{c}_0^\intercal \end{bmatrix} \mathbf{R})$ is statistically close to the distribution of $(\mathbf{A}, \begin{bmatrix} \mathbf{A}' \\ \mathbf{c}_1^\intercal \end{bmatrix})$. □

We now state Theorem 4, which gives a more explicit formulation of Lemma 9. To align the statement with the hybrid argument below, we formulate the experiment in its complementary form, where output 1 denotes extraction failure.

Theorem 4. *Let* SKL-IBE *be the scheme from Construction 2 with security parameter λ and parameters $(q, n, m, \nu, \sigma, \chi, \chi')$. Let $\mathcal{A}$ be any non-uniform* QPT *adversary with polynomial-size advice state τ_λ independent of $(\mathbf{A}, \bar{\mathbf{A}}, \mathbf{B}, \mathbf{y})$, and let $\mathcal{D}$ be any* QPT *distinguisher defined in Figure 2. Assuming the quantum polynomial hardness of* LWE *with polynomial modulus, the* QPT *extractor $\mathcal{E}$ defined in Figure 3 with $\gamma = 1/\mathsf{poly}(\lambda), \delta = 2^{-\Theta(\lambda)}$ satisfies*

$$\Pr\left[1 \leftarrow \mathsf{ExtractExpt}_{\mathcal{A},\mathcal{D},\mathcal{E},\gamma}^{\mathsf{SKL\text{-}IBE}}(1^\lambda)\right] \leq \mathsf{negl}(\lambda),$$

where $\mathsf{ExtractExpt}_{\mathcal{A},\mathcal{D},\mathcal{E},\gamma}^{\mathsf{SKL\text{-}IBE}}(1^\lambda)$ is defined in Figure 4.

Now, we proceed to construct our QPT extractor $\mathcal{E}$ on input $(\mathbf{A}, \mathbf{R}, \hat{\mathbf{y}}, \mathbf{u}, \mathcal{D}, \rho'_{\mathsf{aux}})$ as in Figure 3. By construction, this extractor runs in time $\mathsf{poly}(\lambda, \sigma, 1/\gamma)$.

Proof. We consider the following sequence of hybrids:

$\mathsf{H}_0(\lambda)$**:** This is the experiment defined in Figure 4.

$\underline{\mathcal{E}(\mathbf{A}, \mathbf{R}, \hat{\mathbf{y}}, \mathbf{u}, \mathcal{D}, \rho'_{\mathsf{aux}})}$:

1. Set $\mathbf{x} = \mathbf{0}$, $\varepsilon = \gamma/6$, $\gamma' = \gamma - 3\varepsilon = \gamma/2$.
2. For each $i = 1, 2, \ldots, m$:
   For each $g_i \in [-\sigma\sqrt{m/2}, \ldots, \sigma\sqrt{m/2}]$:
   (a) Let aux store the current state of the quantum distinguisher.
   (b) Run $\overline{\mathsf{ATI}}^{\varepsilon,\delta}_{\mathcal{P}_{\mathcal{D}},(D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}}(i,g_i),D^{\mathbf{A},\mathbf{A}',\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{unif}}),\frac{1}{2}+\gamma'}$ on the residual state in register aux along with enough fresh ancillas initialized to $|0\rangle$.
   (c) Set $\mathbf{x}_i = g_i$ if output is 1. Then, move on to the next guess.
3. Output $\mathbf{x}$.

Fig. 3: The quantum extractor $\mathcal{E}(\mathbf{A}, \mathbf{R}, \hat{\mathbf{y}}, \mathbf{u}, \mathcal{D}, \rho'_{\mathsf{aux}})$.

$\underline{\mathsf{ExtractExpt}^{\mathsf{SKL\text{-}IBE}}_{\mathcal{A},\mathcal{D},\gamma}(1^\lambda)}$:

1. Generate $\mathsf{id}^* \leftarrow \mathcal{A}(1^\lambda)$.
2. Sample $\mathbf{A} \leftarrow\!\$\ \mathbb{Z}_q^{n\times m}$, $(\mathbf{B}, \mathbf{T_B}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$, $\mathbf{R} \leftarrow\!\$\ \{-1,1\}^{m\times m}$, and $\mathbf{u} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma/\sqrt{2}}$. Generate $\left(\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle, \hat{\mathbf{y}}\right) \leftarrow \mathcal{GenGauss}(\mathbf{A}, \sigma)$ for $\mathbf{v} \leftarrow\!\$\ \{0,1\}^m$. Set $\bar{\mathbf{A}} \coloneqq \mathbf{AR} - \mathsf{H}(\mathsf{id}^*)\mathbf{B}$ and $\mathbf{y} \coloneqq \mathbf{ARu} + \hat{\mathbf{y}}$.
3. Simulate $\mathcal{O}_{\mathcal{KG}}$ as follows:
   – For $\mathsf{id}_i = \mathsf{id}^*$, return $\rho_{\mathsf{id}^*} \coloneqq \left(\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle, \hat{\mathbf{y}}, \mathbf{u}\right)$.
   – For $\mathsf{id}_i \neq \mathsf{id}^*$, run $\mathcal{KG}(\cdot)$ with trapdoor $\mathbf{T_B}$ and return ρ_{id_i}.
4. Generate $|\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\mathsf{aux}} \leftarrow \mathcal{A}^{\mathcal{O}_{\mathcal{KG}}}_{\lambda,\mathbf{A},\bar{\mathbf{A}},\mathbf{B},\mathbf{y}}(\bigotimes_{i\in[Q]} \rho_{\mathsf{id}_i} \otimes \tau_\lambda)$ for any $Q = \mathsf{poly}(\lambda)$.
5. Run $b \leftarrow \mathsf{ATI}^{\frac{\gamma}{6},\delta}_{\mathcal{P}_{\mathcal{D}},(D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}},\ D^{\mathbf{A},\mathbf{A}',\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{unif}}),\frac{1}{2}+\gamma}(\rho_{\mathsf{aux}})$, abort and output $\perp$ if $b = 0$.
6. Compute $\mathbf{x} \leftarrow \mathcal{E}(\mathbf{A}, \mathbf{R}, \hat{\mathbf{y}}, \mathbf{u}, \mathcal{D}, \rho'_{\mathsf{aux}})$.
7. Output 1 if $\mathbf{x} \notin \Lambda^{\hat{\mathbf{y}}}_q(\mathbf{A}) \cap \mathcal{B}^m(0, \sigma\sqrt{m/2})$, else output 0.

Fig. 4: The experiment $\mathsf{ExtractExpt}^{\mathsf{SKL\text{-}IBE}}_{\mathcal{A},\mathcal{D},\gamma}(1^\lambda)$.

$\mathsf{H}_1(\lambda)$**:** For step 2, we sample $\mathbf{x} \leftarrow D_{\mathbb{Z}_q^m,\sigma/\sqrt{2}}$ and let $\hat{\mathbf{y}} = \mathbf{Ax}$, then we replace $\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle\!\left\langle\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right|$ with $|\mathbf{x}\rangle\langle\mathbf{x}|$.

$\mathsf{H}_2(\lambda)$**:** For step 5 and inside the extractor algorithm, we replace the distribution $D^{\mathbf{A},\mathbf{A}',\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{unif}}$ with $D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}$.

$\mathsf{H}_3(\lambda)$**:** For step 5, we replace the distribution $D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}}$ with $D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}$.

$\mathsf{H}_4(\lambda)$**:** For the extractor algorithm, we replace the distribution $D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}}(i, g_i)$ with $D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}(i, g_i)$.

Lemma 13. *Assuming the quantum hardness of* LWE*,* H_0 *and* H_1 *are computationally indistinguishable. In other words,*

$$\Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_0}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right] \approx_c \Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_1}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right].$$

Proof. This follows directly from the Gaussian-collapsing property (Lemma 26), assuming the quantum hardness of LWE with the stated parameters:

$$\left(\mathbf{A} \leftarrow\!\$\ \mathbb{Z}_q^{n\times m},\ |\psi_{\mathbf{A},\hat{\mathbf{y}}}\rangle,\ \hat{\mathbf{y}} \in \mathbb{Z}_q^n\right) \approx_c \left(\mathbf{A} \leftarrow\!\$\ \mathbb{Z}_q^{n\times m},\ |\mathbf{x}'\rangle,\ \mathbf{A}\cdot\mathbf{x}' \in \mathbb{Z}_q^n\right),$$

where $(|\psi_{\mathbf{A},\hat{\mathbf{y}}}\rangle, \hat{\mathbf{y}}) \leftarrow \mathit{GenGauss}(\mathbf{A}, \sigma)$ and where $\mathbf{x}' \leftarrow D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and the fact that the Pauli-Z operator $\mathbf{Z}_q^{\lfloor \frac{q}{\nu} \rfloor \cdot \mathbf{v}}$ has no effect on the collapsed state. □

Lemma 14. H_1 *and* H_2 *are statistically indistinguishable. In other words,*

$$\Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_1}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right] \approx_s \Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_2}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right].$$

Proof. The statistical indistinguishability follows directly from Lemma 12. □

Lemma 15. *Assuming the quantum hardness of* LWE, H_2 *and* H_3 *are computationally indistinguishable. In other words,*

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_2}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_3}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right]\right| \leq \mathsf{negl}(\lambda).$$

Proof. Assuming there exists a QPT adversary that can distinguish the two hybrids with non-negligible advantage $\frac{1}{\mu}$ for some polynomial $\mu(\lambda)$, we construct a reduction $\mathcal{R}$ to break the quantum hardness of decisional $\mathsf{LWE}^m_{n,q,\alpha q}$ with non-negligible advantage. By standard averaging argument, for at least $\frac{1}{2\mu}$ fraction of samples $(\mathbf{A}, \mathbf{x}')$, our adversary can distinguish the two hybrids with advantage at least $\frac{1}{2\mu}$. For such samples, by Lemma 24, there exists a quantum circuit $\mathcal{C}$ of size $\mathsf{poly}(\lambda, 1/\varepsilon, \log(1/\delta))$ such that $\mathcal{C}(\mathbf{A}, \mathbf{x}', \cdot)$ succeed in distinguishing samples from $(D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}}, D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}})$ and $(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}, D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}})$ with advantage at least $\frac{1}{8\mu^3\mathsf{poly}(\lambda,1/\varepsilon,\log(1/\delta))}$, which is also inverse polynomial and can be written as $\frac{1}{\mu'}$.

From Lemma 10, we have that $D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}} \approx_s \bar{D}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{lwe}}$. Combined with the fact that $D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}$ can be sampled independently of $\bar{D}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{lwe}}$ and $D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}$ given $(\mathbf{R}, \mathbf{u})$, $\mathcal{C}(\mathbf{A}, \mathbf{x}', \cdot)$ can distinguish samples from $(\bar{D}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{lwe}}, D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}})$ and $(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}, D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}})$ with advantage at least $\frac{1}{\mu'} - \mathsf{negl}(\lambda)$. By definition, the only difference between $\bar{D}^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{lwe}}$ and $D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}$ is from the vector $\mathbf{c}_0$, which is either an LWE sample or a uniformly random vector in $\mathbb{Z}_q^m$. Thus, the reduction $\mathcal{R}$ can simulate the view of our adversary and distinguish the LWE challenge with advantage at least $\frac{1}{\mu'} - \mathsf{negl}(\lambda)$. This contradicts the quantum hardness of decisional $\mathsf{LWE}^m_{n,q,\chi}$, thus the proof is complete. □

Lemma 16. *Assuming the quantum hardness of* LWE, H_3 *and* H_4 *are computationally indistinguishable. In other words,*

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_3}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_4}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right]\right| \leq \mathsf{negl}(\lambda).$$

Proof. We define intermediate hybrids $\mathsf{H}_{3,k}$ for $0 \leq k \leq \sigma\sqrt{2m^3} - 1$ as follows:

1-5. Same as H_3.
6.1. Set $\mathbf{x} = \mathbf{0}$, $\varepsilon = \gamma/6$, $\gamma' = \gamma - 3\varepsilon = \gamma/2$, $\mathsf{cnt} = 0$.
6.2. For each $i \in [m]$:
For each $g_i \in [-\sigma\sqrt{m/2}, \ldots, \sigma\sqrt{m/2}]$:
(a) Let aux store the current state of $\mathcal{D}$ and set $\mathsf{cnt} := \mathsf{cnt} + 1$.

(b) If $\mathsf{cnt} \leq k$, run $\overline{\mathsf{ATI}}^{\varepsilon,\delta}_{\mathcal{P}_\mathcal{D},(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}(i,g_i),D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}),1/2+\gamma'}$ on the residual state in register aux along with enough fresh ancillas initialized to $|0\rangle$, otherwise run $\overline{\mathsf{ATI}}^{\varepsilon,\delta}_{\mathcal{P}_\mathcal{D},(D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}}(i,g_i),D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}),1/2+\gamma'}$ as in H_3.

(c) Set $\mathbf{x}_i = g_i$ if output is 1. Then, move on to the next guess.

6.3. Output $\mathbf{x}$.

7. Same as H_3.

We observe that $\mathsf{H}_{3,0}$ is identical to H_3 and $\mathsf{H}_{3,\sigma\sqrt{2m^3}}$ is identical to H_4. Therefore, it suffices to prove that for $0 \leq k \leq \sigma\sqrt{2m^3} - 1$:

$$\left|\Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_{3,k}}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right] - \Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_{3,k+1}}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right]\right| \leq \mathsf{negl}(\lambda)/\sigma. \tag{4}$$

The proof for Equation (4) is almost the same as the proof of Lemma 14, except that we apply Lemma 25 for $\overline{\mathsf{ATI}}$ instead of Lemma 24 for ATI, the parameter μ is set as $1/\mathsf{poly}(\lambda,\sigma)$ instead of $1/\mathsf{poly}(\lambda)$. □

Lemma 17. H_4 *outputs* 1 *with negligible probability. In other words,*

$$\Pr\left[1 \leftarrow \mathsf{Expt}^{\mathsf{H}_4}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right] \leq \mathsf{negl}(\lambda).$$

Proof. First, we make a modification to H_4, where it will output 1 if the extracted $\mathbf{x} \neq \mathbf{x}'$, which is implied by H_4 outputs 1 (that is, when $\mathbf{x} \notin \Lambda^{\hat{\mathbf{y}}}_q(\mathbf{A}) \cap \mathcal{B}^m(0,\sigma\sqrt{m/2})$ holds). To prove Lemma 17, it suffices to prove the experiment in Figure 5 outputs 1 with negligible probability.

At its core, the extractor algorithm applies a sequence of projective measurements $\overline{\mathsf{ATI}}$ on system aux. With the Quantum Union Bound lemma, we can argue the probability of this experiment, or this sequence of measurements, outputs 1 by examining the outcome probability of each measurement. To this end, we define a sub-experiment in Figure 6, which concerns a single projective measurement.

By Lemma 1, we have the following:

$$\Pr\left[1 \leftarrow \mathsf{ExtractorExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda)\right] \leq 4\sum_{i=1}^{m}\sum_{g_i=-\sigma\sqrt{m/2}}^{\sigma\sqrt{m/2}} \Pr\left[1 \leftarrow \mathsf{ExtSubExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda,i,g_i,\mathbf{x}')\right].$$

Now, if we can show the probability of the sub-experiment is always at most negligible for any input $(i,g_i,\mathbf{x}')$, then our proof is complete. This can be done by considering two cases: whether the guess g_i is correct or not.

Correct guess ($g_i = x'_i$): Let ρ'_{aux} be the residual state in register aux after running $\overline{\mathsf{ATI}}^{\varepsilon,\delta}_{\mathcal{P}_\mathcal{D},(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}},\, D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}),\frac{1}{2}+\gamma}$ on the initial state ρ_{aux} and getting output $b = 1$. From Lemma 23, running $\overline{\mathsf{ATI}}^{\varepsilon,\delta}_{\mathcal{P}_\mathcal{D},(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}},\, D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}),\frac{1}{2}+\gamma-3\varepsilon}$ on ρ'_{aux} will output 1 with probability at least $1-3\delta$. As our sub-experiment will output 1 only when the corresponding $\overline{\mathsf{ATI}}$ outputs 0, we have the following:

$$\Pr\left[1 \leftarrow \mathsf{ExtSubExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda,i,g_i,\mathbf{x}')\right] \leq 3\delta.$$

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

1. Generate $\mathsf{id}^* \leftarrow \mathcal{A}(1^\lambda)$.
2. Sample $\mathbf{A} \leftarrow\$ \mathbb{Z}_q^{n\times m}$, $(\mathbf{B}, \mathbf{T_B}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$, $\mathbf{R} \leftarrow\$ \{-1,1\}^{m\times m}$, and $\mathbf{u} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma/\sqrt{2}}$. Generate $\left(\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle, \hat{\mathbf{y}}\right) \leftarrow \mathit{GenGauss}(\mathbf{A},\sigma)$ for $\mathbf{v} \leftarrow\$ \{0,1\}^m$. Set $\bar{\mathbf{A}} := \mathbf{AR} - \mathsf{H}(\mathsf{id}^*)\mathbf{B}$ and $\mathbf{y} := \mathbf{ARu} + \hat{\mathbf{y}}$.
3. Simulate $\mathcal{O}_{\mathcal{KG}}$ as follows:
   - For $\mathsf{id}_i = \mathsf{id}^*$, return $\rho_{\mathsf{id}^*} := \left(\left|\psi^{\mathbf{v}}_{\mathbf{A},\hat{\mathbf{y}}}\right\rangle, \hat{\mathbf{y}}, \mathbf{u}\right)$.
   - For $\mathsf{id}_i \neq \mathsf{id}^*$, run $\mathcal{KG}(\cdot)$ with $\mathbf{T_B}$ and return ρ_{id_i}.
4. Generate $|\mathbf{w}\rangle\langle\mathbf{w}| \otimes \rho_{\mathsf{aux}} \leftarrow \mathcal{A}^{\mathcal{O}_{\mathcal{KG}}}_{\lambda,\mathbf{A},\bar{\mathbf{A}},\mathbf{B},\mathbf{y}}(\bigotimes_{i\in[Q]} \rho_{\mathsf{id}_i} \otimes \tau_\lambda)$ for any $Q = \mathsf{poly}(\lambda)$.
5. Run $b \leftarrow \mathsf{ATI}^{\frac{\gamma}{6},\delta}_{\mathcal{P}_\mathcal{D},(D^{\mathbf{A},\mathbf{R},\hat{\mathbf{y}},\mathbf{y}}_{\mathsf{lwe}},\ D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}),\frac{1}{2}+\gamma}(\rho_{\mathsf{aux}})$. Continue if $b = 1$, otherwise abort and output $\perp$.
6. Set $\mathbf{x} = \mathbf{0}$, $\varepsilon = \gamma/6$, $\gamma' = \gamma - 3\varepsilon = \gamma/2$.
7. For each $i \in [m]$:
   For each $g_i \in [-\sigma\sqrt{m/2}, \ldots, \sigma\sqrt{m/2}]$:
   (a) Let aux store the current state of the quantum distinguisher.
   (b) Run $\overline{\mathsf{ATI}}^{\varepsilon,\delta}_{\mathcal{P}_\mathcal{D},(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}(i,g_i),D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}),1/2+\gamma'}$ on the residual state in register aux along with enough fresh ancillas initialized to $|0\rangle$.
   (c) Set $\mathbf{x}_i = g_i$ if output is 1. Then, move on to the next guess.
8. Output 1 if $\mathbf{x} \neq \mathbf{x}'$, otherwise output 0.

Fig. 5: Experiment of $\mathsf{ExtractorExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda)$.

$\underline{\mathsf{ExtSubExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')}$:

1-6. Same as in $\mathsf{ExtractorExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda)$.
7. Run $\overline{\mathsf{ATI}}^{\varepsilon,\delta}_{\mathcal{P}_\mathcal{D},(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}(i,g_i),D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}),\frac{1}{2}+\gamma'}$ on the residual state in register aux along with enough fresh ancillas initialized to $|0\rangle$.
8. Output 1 if $(g_i \neq x_i \wedge 1 \leftarrow \overline{\mathsf{ATI}})$ or $(g_i = x_i \wedge 0 \leftarrow \overline{\mathsf{ATI}})$, otherwise output 0.

Fig. 6: Experiment of $\mathsf{ExtSubExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')$.

Incorrect guess ($g_i \neq x'_i$): From Lemma 11, we have $D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{cg}}(i,g_i) = D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}$, thus we have

$$\Pr\left[1 \leftarrow \mathsf{ExtSubExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')\right] = \mathrm{Tr}\left[\overline{\mathsf{ATI}}^{\varepsilon,\delta}_{\mathcal{P}^\mathcal{D},(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}},\ D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}),\frac{1}{2}+\gamma'} \rho'_{\mathsf{aux}}\right].$$

As $\mathcal{P}^\mathcal{D}_{(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}},\ D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}})}$ only has eigenvalue $\frac{1}{2} < \frac{1}{2}+\gamma'-\varepsilon$ (any distinguisher cannot do better than guessing a random bit), running $\overline{\mathsf{ATI}}^{\varepsilon,\delta}_{\mathcal{P}^\mathcal{D},(D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}},\ D^{\mathbf{A},\mathbf{R},\mathbf{x},\mathbf{u}}_{\mathsf{wg}}),\frac{1}{2}+\gamma'}$ on any state will output 1 with probability at most δ,

$$\Pr\left[1 \leftarrow \mathsf{ExtSubExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')\right] \leq \delta.$$

Combining the two cases with δ defined to be exponentially small $\delta \leq \mathsf{negl}(\lambda)/\sigma$, we have bounded the experiment in Figure 5 to output 1 with $\mathsf{negl}(\lambda)/\sigma$:

$$\Pr\left[1 \leftarrow \mathsf{ExtSubExpt}_{\mathcal{A},\mathcal{D}}(1^\lambda, i, g_i, \mathbf{x}')\right] \leq \mathsf{negl}(\lambda)/\sigma. \qquad \square$$

With Lemmas 13 to 17, the proof for Theorem 4 is complete. $\square$

Acknowledgments. We thank the anonymous SAC reviewers for suggestions for improving the presentation of the paper. H. N. Pham was supported by the Program QuanTEdu-France (ANR-22-CMAS-0001) France 2030. D. H. Phan was supported in part by the France 2030 ANR Project ANR-22-PECY-003 SecureCompute. Q.-H. Vu was supported in part by the French ANR project CTQC (ANR-25-CE48-7926-01 CTQC). W. Wen was supported in part by the HELO ANR grant (ANR-24-CE48-4293).

A Preliminaries

A.1 Lattices Background

A m-dimensional *lattice* $\Lambda \subset \mathbb{R}^m$ is a discrete subgroup of $\mathbb{R}^m$. Given positive integers $n, m, q \in \mathbb{N}$ and a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, let $\Lambda_q^\perp(\mathbf{A})$ denote a *q-ary* lattice generated by $\mathbf{A}$: $\{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\mathbf{x} = \mathbf{0} \pmod q\}$. For $\mathbf{y} \in \mathbb{Z}_q^n$, we denote the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ given by $\{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\mathbf{x} = \mathbf{y} \pmod q\}$. For a lattice Λ and positive real $\varepsilon > 0$, the *smoothing parameter* $\eta_\varepsilon(\Lambda)$ of Λ is defined to be the smallest positive s such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$, where Λ^* is the dual lattice of Λ.

Discrete Gaussians. Let $\Sigma \in \mathbb{R}^{m\times m}$ be a full-rank, symmetric, and positive definite, let $\mathbf{c} \in \mathbb{R}^m$. The Gaussian function with center $\mathbf{c}$ and covariance matrix $\sqrt{\Sigma}$ is defined by the probability distribution function $\rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbf{x}) = \exp\big(-\pi \cdot (\mathbf{x}-\mathbf{c})^\top \Sigma^{-1}(\mathbf{x}-\mathbf{c})\big)$, for any $\mathbf{x} \in \mathbb{R}^m$. In the spherical case $\Sigma = \sigma^2\mathbf{I}$, Gaussian function can be simplified as $\rho_\sigma(\mathbf{x}) = \exp\big(-\pi \cdot \|\mathbf{x}-\mathbf{c}\|^2/\sigma^2\big)$. When $\mathbf{c} = \mathbf{0}$, we omit it from the subscript.

Definition 4 (Discrete Gaussian Distribution). *For a full-rank, symmetric, positive definite $m \times m$ matrix Σ, we define the m-dimension discrete Gaussian distribution over the lattice $\mathbb{Z}^m$, $\mathcal{D}_{\mathbb{Z}^m,\sqrt{\Sigma},\mathbf{c}}$ with center $\mathbf{c} \in \mathbb{R}^m$ and standard deviation parameter matrix $\sqrt{\Sigma}$ by*

$$\forall \mathbf{x} \in \mathbb{Z}^m : \ \mathcal{D}_{\mathbb{Z}^m,\sqrt{\Sigma},\mathbf{c}}(\mathbf{x}) = \rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbf{x})/\rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbb{Z}^m),$$

where $\rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbb{Z}^m) = \sum_{\mathbf{x}\in\mathbb{Z}^m} \rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbf{x})$.

Definition 5 (Continuous Gaussian Distribution). *Let $\mathbb{R}^m$ be the Euclidean space, and let Σ be a positive definite matrix in $\mathbb{R}^{m\times m}$. The Gaussian (or normal) distribution $\mathcal{N}_{m,\sqrt{\Sigma},\mathbf{c}}$ over $\mathbb{R}^m$ with center $\mathbf{c} \in \mathbb{R}^m$ and covariance matrix Σ is defined by the density function:*

$$p(\mathbf{x}, \Sigma, \mathbf{c}) = \frac{1}{\det(\Sigma)^{1/2}}\, \rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbf{x}).$$

Lemma 18 (Generalized Leftover Hash Lemma [ABB10, Lemma 13]). *Let n, m, q be integers such that $m \geq 2n \log q$ and $q > 2$ is prime. Then, for all fixed vectors $e \in \mathbb{Z}_q^m$ and all $k = \mathsf{poly}(n)$, the statistical distance between the following distributions is $\mathsf{negl}(n)$:*

- $\left(\mathbf{A}, \mathbf{AR}, \mathbf{e}^\intercal\mathbf{R}\right)$ *where* $\mathbf{A} \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times m}$ *and* $\mathbf{R} \leftarrow\!\!\$\ \{-1,1\}^{m\times k}$.
- $\left(\mathbf{A}, \mathbf{U}, \mathbf{e}^\intercal\mathbf{R}\right)$ *where* $\mathbf{A} \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times m}$, $\mathbf{U} \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times k}$, *and* $\mathbf{R} \leftarrow\!\!\$\ \{-1,1\}^{m\times k}$.

Lemma 19 ([Reg05, Claim 3.9]). *Let Λ be a lattice, let $\mathbf{u} \in \mathbb{R}^m$ be any vector, let $r, s > 0$ be two real numbers, and let $t := \sqrt{r^2 + s^2}$. Assume that $rs/t \geq \eta_\varepsilon(\Lambda)$ for some $\varepsilon < \frac{1}{2}$. Consider the continuous distribution Y on $\mathbb{R}^m$ obtained by sampling from $\mathcal{D}_{\Lambda+\mathbf{u},r}$ and then adding a noise vector taken from $\mathcal{N}_{m,s}$. Then the statistical distance between Y and $\mathcal{N}_{m,t}$ is at most 4ε.*

Lemma 20 ([Pei10, Theorem 3.1]). *Let $\Sigma_1, \Sigma_2 > 0$ be two $m \times m$ positive definite matrices, with $\Sigma = \Sigma_1 + \Sigma_2 > 0$ and $\Sigma_3^{-1} = \Sigma_1^{-1} + \Sigma_2^{-1} > 0$. Let Λ_1 be a lattice satisfying $\lambda_{\sqrt{\Sigma_1}}^{\max} \geq \eta_\varepsilon(\Lambda_1)$ for some $\varepsilon \leq \frac{1}{2}$. Let $\mathbf{v}_1 \in \mathbb{R}^m$ be arbitrary. Consider the following probabilistic experiment: choose $\mathbf{x}_2 \leftarrow \mathcal{N}_{m,\sqrt{\Sigma_2}}$, then choose $\mathbf{x}_1 \leftarrow \mathcal{D}_{\Lambda_1+\mathbf{v}_1-\mathbf{x}_2,\sqrt{\Sigma_1}}$. The marginal distribution of $\mathbf{x}_1 + \mathbf{x}_2$ is within statistical distance 8ε of $\mathcal{D}_{\Lambda_1+\mathbf{v}_1,\sqrt{\Sigma}}$.*

Lemma 21 ([BD20, Proposition 3.2]). *Let $\mathbf{F} \in \mathbb{Z}^{m\times n}$ be an arbitrary matrix with spectral norm $r_{\mathbf{F}}$. Let $r, r_1 > 0$ be such that $r > r_1 \cdot r_{\mathbf{F}}$. Let $\mathbf{e}_1 \leftarrow \mathcal{N}_{n,r_1}$ and let $\mathbf{e}_2 \leftarrow \mathcal{N}_{m,\sqrt{\Sigma}}$ for $\Sigma = r^2 I_{m\times m} - r_1^2\mathbf{F}\mathbf{F}^\top$. Then the random variable $\mathbf{e} = \mathbf{F}\mathbf{e}_1 + \mathbf{e}_2$ is distributed according to $\mathcal{N}_{m,r}$.*

Definition 6 (Short Integer Solution, [Ajt96]). *Let $n, m \in \mathbb{N}$, $q \geq 2$ be a modulus and let $\beta > 0$ be a parameter. The Short Integer Solution $\mathsf{SIS}_{n,q,\beta}^m$ problem is to find a short solution $\mathbf{x} \in \mathbb{Z}^m$ with $\|\mathbf{x}\| \leq \beta$ such that $\mathbf{Ax} = \mathbf{0} \pmod q$ given as input a matrix $\mathbf{A} \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times m}$.*

Definition 7 (Learning with Errors, [Reg05]). *Let $n, m \in \mathbb{N}$ be integers, let $q \geq 2$ be a modulus and let χ be the noise bound. The Learning with Errors $\mathsf{LWE}_{n,q,\chi}^m$ problem is to distinguish between the following samples:*

$$\left\{\left(\mathbf{A}, \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal\right) \ : \ \begin{array}{c}\mathbf{A} \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times m},\\ \mathbf{s} \leftarrow\!\!\$\ \mathbf{Z}_q^n,\ \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\chi}\end{array}\right\} \text{ and } \left\{\left(\mathbf{A}, \mathbf{u}^\intercal\right) \ : \ \begin{array}{c}\mathbf{A} \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times m},\\ \mathbf{u} \leftarrow\!\!\$\ \mathbb{Z}_q^m\end{array}\right\}$$

In this work, we rely on the quantum hardness of $\mathsf{LWE}_{n,q,\chi}^m$ assumption against quantum adversaries for $m = \Omega(n \log q)$, in a parameter regime where it implies $\mathsf{SIS}_{n,q,\beta}^m$ [SSTX09].

Trapdoors for lattices. We use the following properties of short basis trapdoors for the q-ary lattice $\Lambda_q^\perp(\mathbf{A})$. For $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and $\mathbf{u} \in \mathbb{Z}_q^n$, we write $\mathbf{A}_\sigma^{-1}(\mathbf{u})$ to denote a sample $\mathbf{x} \leftarrow D_{\Lambda_q^{\mathbf{u}}(\mathbf{A}),\sigma}$.

Lemma 22 ([GPV08, AP09, MP12]). *Let $n, m \in \mathbb{N}$ and $q \in \mathbb{N}$ be a prime with $m = \Omega(n \log q)$. There exists the following randomized algorithms:*

– $\mathsf{GenTrap}(1^n, 1^m, q)$*: on input* 1^n, 1^m *and* q*, returns a matrix* $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ *and a trapdoor* $\mathbf{T_A}$ *such that the distribution of* $\mathbf{A}$ *is negligibly close to uniform and* $\mathbf{T_A}$ *is a* $(20n\log q)$*-good lattice trapdoor with all but negligible probability.*
– $\mathsf{ExtendRight}(\mathbf{A}, \mathbf{B}, \mathbf{T_A})$([CHKP10])*: on input full-rank matrices* $\mathbf{A}, \mathbf{B} \in \mathbb{Z}_q^{n\times m}$ *and a trapdoor* $\mathbf{T_A}$ *of* $\Lambda_q^\perp(\mathbf{A})$*, outputs a trapdoor* $\mathbf{T}_{[\mathbf{A}|\mathbf{B}]}$ *of* $\Lambda_q^\perp([\mathbf{A} \mid \mathbf{B}])$ *such that* $\|\mathbf{T_A}\|_{\mathrm{GS}} = \left\|\mathbf{T}_{[\mathbf{A}|\mathbf{B}]}\right\|_{\mathrm{GS}}$.
– $\mathsf{ExtendLeft}(\mathbf{A}, \mathbf{B}, \mathbf{T_B}, \mathbf{R})$([ABB10])*: on input full-rank matrices* $\mathbf{A}, \mathbf{B} \in \mathbb{Z}_q^{n\times m}$ *and a trapdoor* $\mathbf{T_B}$ *of* $\Lambda_q^\perp(\mathbf{B})$*, outputs a trapdoor* $\mathbf{T}_{[\mathbf{A}|\mathbf{AR}+\mathbf{B}]}$ *of* $\Lambda_q^\perp([\mathbf{A} \mid \mathbf{B}])$ *such that* $\left\|\mathbf{T}_{[\mathbf{A}|\mathbf{AR}+\mathbf{B}]}\right\|_{\mathrm{GS}} \le \|\mathbf{T_B}\|_{\mathrm{GS}} \cdot \left(1 + \|\mathbf{R}\|_2\right)$.

A.2 Cryptographic Primitives

Definition 8 (Pseudorandom Function). *Let* $\{F_K : \{0,1\}^{\ell_1} \to \{0,1\}^{\ell_2} \mid K \in \{0,1\}^\lambda\}$ *be a family of polynomially computable functions, where* ℓ_1 *and* ℓ_2 *are some polynomials of* λ*. We say that* F *is a post-quantum secure pseudorandom function (PRF) family if, for any* QPT *distinguisher* $\mathcal{A}$*, it holds that*

$$\left|\Pr\left[\mathcal{A}^{F_K(\cdot)}(1^\lambda) = 1 \mid K \leftarrow \{0,1\}^\lambda\right] - \Pr\left[\mathcal{A}^{R(\cdot)}(1^\lambda) = 1 \mid R \leftarrow \mathcal{U}\right]\right| \le \mathsf{negl}(\lambda),$$

where $\mathcal{U}$ *is the set of all functions from* $\{0,1\}^{\ell_1}$ *to* $\{0,1\}^{\ell_2}$.

Theorem 5 ([GGM84, Ajt96]). *Assuming the quantum hardness of* $\mathsf{SIS}^m_{n,q,\beta}$*, there exists a post-quantum secure PRF.*

A.3 Threshold Implementation

In this section, we recall technical lemmas underlying Zhandry's measurement procedure [Zha20, ALL+21], which provides a notion of *testing* adversarial quantum states in security games. Unlike the classical setting, where the challenger checks returned classical strings, standard classical notions of "testing" for quantum states can fail in several ways [Zha20]. Zhandry's procedure addresses this by estimating the state's success probability while causing only limited disturbance, enabling the repeated tests needed in quantum reductions. We refer the reader to [Zha20, ALL+21] for a detailed treatment of this notion.

Theorem 6 (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* $(\mathsf{TI}_\gamma(P), \mathbb{I} - \mathsf{TI}_\gamma(P))$ *such that*

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

While the threshold implementation defined above can not be efficiently in general, an approximate version of the threshold implementation can be implemented efficiently as long as the POVM is a mixture of projective measurements.

Lemma 23 (Approximate Threshold Implementation (ATI), [ALL+21, Corollary 1]). *Let $\mathcal{P}_D = (P_D, Q_D)$ be a binary outcome POVM over Hilbert space $\mathcal{H}$ that is a mixture of projective measurements over some distribution D. Let $\varepsilon, \delta, \gamma \in (0,1)$. Then, there exists an efficient binary-outcome quantum algorithm $\mathsf{ATI}^{\varepsilon,\delta}_{\mathcal{P}_D,\gamma}$, interpreted as the POVM element corresponding to outcome 1, such that the following holds:*

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

To apply hybrid arguments between $\mathsf{ATI}^{\varepsilon,\delta}_{\mathcal{P},D_0,\gamma}$ and $\mathsf{ATI}^{\varepsilon,\delta}_{\mathcal{P},D_1,\gamma}$, we will use the following results from [AHH24] that allow us to preserve the distinguishing advantage γ. Moreover, since ATI is not a projector and may essentially change the state even when it outputs 1 with overwhelming probability, they introduced the notion of projective ATI, denoted $\overline{\mathsf{ATI}}$, to minimize the disturbance on the state by performing uncomputation.

Lemma 24 (Post-ATI Distinguisher, [AHH24, Lemma 5]). *Let $\mathcal{P}$ be a collection of projective measurements indexed by some set $\mathcal{I}$, implemented by a quantum circuit of size $|\mathcal{P}|$. Let D_0, D_1 be two efficiently samplable distributions over $\mathcal{I}$. For any polynomial μ, any quantum state ρ and any (possibly quantum) predicate $h : \{0,1\} \times \mathcal{D}(\mathcal{H}) \to \{0,1\}$ with circuit size $|h|$, if*

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

there exists a quantum circuit $\mathcal{C}$ of size $\mathsf{poly}\left(\lambda, \frac{1}{\varepsilon}, \log\left(\frac{1}{\delta}\right), \mu, |\mathcal{P}|, |h|\right)$, such that

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

Lemma 25 (Post-$\overline{\mathsf{ATI}}$ Distinguisher, [AHH24, Lemma 6]). *Let $\mathcal{P}$ be a collection of projective measurements indexed by some set $\mathcal{I}$, implemented by a quantum circuit of size $|\mathcal{P}|$. Let D_0, D_1 be two efficiently samplable distributions over $\mathcal{I}$. For any polynomial μ, any quantum state ρ and any (possibly quantum) predicate $h : \{0,1\} \times \mathcal{D}(\mathcal{H}_2) \to \{0,1\}$ with circuit size $|h|$, if*

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

there exists a quantum circuit $\mathcal{C}$ of size $\mathsf{poly}\left(\lambda, \frac{1}{\varepsilon}, \log\left(\frac{1}{\delta}\right), \mu, |\mathcal{P}|, |h|\right)$, such that

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

A.4 Gaussian Coset States

Our analysis relies on the computational intractability of distinguishing whether a Gaussian coset state has been measured in the computational basis. This property was proved under subexponential modulus-to-noise ratio in [LZ19, Por23], and later extended to any polynomial modulus-to-noise ratio in [LMZ23].

Lemma 26 (Gaussian-collapsing property, [LMZ23, Theorem 5, adapted]). *Let $n \in \mathbb{N}$ and q parameterized by $\lambda \in \mathbb{N}$. Then, the following samples are computationally indistinguishable assuming the quantum hardness of decisional* LWE*:*

$$\left(\mathbf{A} \leftarrow\!\$\, \mathbb{Z}_q^{n\times m},\ |\psi_{\mathbf{A},\mathbf{y}}\rangle,\ \mathbf{y} \in \mathbb{Z}_q^n\right) \approx_c \left(\mathbf{A} \leftarrow\!\$\, \mathbb{Z}_q^{n\times m},\ |\mathbf{x}_0\rangle,\ \mathbf{A}\mathbf{x}_0 \in \mathbb{Z}_q^n\right)$$

where $(|\psi_{\mathbf{A},\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathit{GenGauss}(\mathbf{A}, \sigma)$ *and* $\mathbf{x}_0 \leftarrow D_{\mathbb{Z}_q^m, \sigma/\sqrt{2}}$.

For classical revocation, we rely on the *certified-everlasting collapsing* property of Gaussian coset states, conjectured in [Por23] and proved in [BKP23] under LWE with subexponential modulus. The proof reduces to the existence of any *collapsing* and *collision-resistant* hash function, which can be instantiated by the SIS hash function [Ajt96]. By our parameter choice, the required SIS collision resistance follows from the LWE assumption. Using Lemma 26, we obtain the following theorem.

Theorem 7 (Certified-everlasting target-collapsing property of Gaussian coset states, [BKP23]). *Let $n \in \mathbb{N}$ and q be a prime modulus with $m \geq 2n \log q$, each parameterized by security parameter $\lambda \in \mathbb{N}$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and $\chi > 2\sqrt{n}$. Assuming the hardness of the* $\mathsf{LWE}^m_{n,q,\chi}$*, it holds for any adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ with a* QPT *algorithm $\mathcal{A}_0$ and an unbounded $\mathcal{A}_1$:*

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

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

1. *$\mathcal{Ch}$ generates a Gaussian coset state $(|\psi_{\mathbf{A},\mathbf{y}}\rangle, \mathbf{y}) \in \mathbb{Z}_q^n \leftarrow \mathit{GenGauss}(\mathbf{A}, \sigma)$.*
2. *If $b = 0$, $\mathcal{Ch}$ does nothing. Else, $\mathcal{Ch}$ measures the state in the computational basis. Next, $\mathcal{Ch}$ sends the resulting state together with $\mathbf{y}$ to $\mathcal{A}_0$.*
3. *$\mathcal{A}_0$ sends a classical certificate $\mathbf{w} \in \mathbb{Z}_q^m$ to $\mathcal{Ch}$ and initializes $\mathcal{A}_1$ with its residual (internal) state.*
4. *$\mathcal{Ch}$ checks if $\mathbf{w}$ satisfies $\mathbf{A}\mathbf{w} = \mathbf{y} \pmod q$ and $\|\mathbf{w}\| \leq \sigma\sqrt{m/2}$. If true, $\mathcal{A}_1$ is run until it outputs b', otherwise $b' \leftarrow\!\$\, \{0,1\}$. The final output is b'.*

B Missing Proofs

Claim B.1. *Suppose there are three probabilistic events A, B, C, if $\Pr[A \mid B] \geq 1 - \mathsf{negl}(\lambda)$ and $\Pr[B \wedge C] \geq 1/\mathsf{poly}(\lambda)$, then $\Pr[A \wedge C] \geq 1/\mathsf{poly}(\lambda) - \mathsf{negl}(\lambda)$.*

Proof. Since $\Pr[A \mid B] \geq 1 - \mathsf{negl}(\lambda)$, we have $\Pr[A \wedge B] = \Pr[B] \cdot \Pr[A \mid B] \geq \Pr[B] - \mathsf{negl}(\lambda)$. Therefore, we have $\Pr\left[B \wedge \overline{A}\right] \leq \mathsf{negl}(\lambda)$. Then we can deduce:

$$\begin{aligned}\Pr[A \wedge C] &\geq \Pr[A \wedge B \wedge C] = \Pr[B \wedge C] - \Pr\left[B \wedge C \wedge \overline{A}\right] \\ &\geq 1/\mathsf{poly}(\lambda) - \Pr\left[B \wedge \overline{A}\right] \geq 1/\mathsf{poly}(\lambda) - \mathsf{negl}(\lambda).\end{aligned}$$

□

References

ABB10. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 553–572. Springer, Berlin, Heidelberg, May / June 2010. 5, 11, 12, 24, 25

AHH24. Prabhanjan Ananth, Zihan Hu, and Zikuan Huang. Quantum key-revocable dual-regev encryption, revisited. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part III*, volume 15366 of *LNCS*, pages 257–288. Springer, Cham, December 2024. 2, 3, 4, 6, 10, 11, 16, 17, 26

Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996. 24, 25, 27

AKN+23. Shweta Agrawal, Fuyuki Kitagawa, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa. Public key encryption with secure key leasing. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 581–610. Springer, Cham, April 2023. 1, 2, 3, 4, 6, 8

ALL+21. Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. New approaches for quantum copy-protection. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 526–555, Virtual Event, August 2021. Springer, Cham. 25, 26

AP09. Joël Alwen and Chris Peikert. Generating Shorter Bases for Hard Random Lattices. In Susanne Albers and Jean-Yves Marion, editors, *Proceedings of the 26th Annual Symposium on the Theoretical Aspects of Computer Science*, Proceedings of the 26th Annual Symposium on the Theoretical Aspects of Computer Science, pages 75–86, Freiburg, Germany, February 2009. IBFI Schloss Dagstuhl. 24

APV23. Prabhanjan Ananth, Alexander Poremba, and Vinod Vaikuntanathan. Revocable cryptography from learning with errors. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part IV*, volume 14372 of *LNCS*, pages 93–122. Springer, Cham, November / December 2023. 1, 2, 3, 4, 9, 10, 13

BCM+18. Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In Mikkel Thorup, editor, *59th FOCS*, pages 320–331. IEEE Computer Society Press, October 2018. 6

BD20. Zvika Brakerski and Nico Döttling. Hardness of LWE on general entropic distributions. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 551–575. Springer, Cham, May 2020. 3, 4, 24

BGK08. Alexandra Boldyreva, Vipul Goyal, and Virendra Kumar. Identity-based encryption with efficient revocation. In Peng Ning, Paul F. Syverson, and Somesh Jha, editors, *ACM CCS 2008*, pages 417–426. ACM Press, October 2008. 2

BGK+24. James Bartusek, Vipul Goyal, Dakshita Khurana, Giulio Malavolta, Justin Raizes, and Bhaskar Roberts. Software with certified deletion. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part IV*, volume 14654 of *LNCS*, pages 85–111. Springer, Cham, May 2024. 3

BKP23. James Bartusek, Dakshita Khurana, and Alexander Poremba. Publicly-verifiable deletion via target-collapsing functions. In Helena Handschuh and

Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 99–128. Springer, Cham, August 2023. 27

BKVV20. Zvika Brakerski, Venkata Koppula, Umesh Vazirani, and Thomas Vidick. Simpler Proofs of Quantumness. *LIPIcs, Volume 158, TQC 2020*, 158, 2020. 7

CGJL25. Orestis Chardouvelis, Vipul Goyal, Aayush Jain, and Jiahui Liu. Quantum key leasing for PKE and FHE with a classical lessor. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 248–277. Springer, Cham, May 2025. 2, 3, 4, 6

CHKP10. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 523–552. Springer, Berlin, Heidelberg, May / June 2010. 5, 25

CHV23. Céline Chevalier, Paul Hermouet, and Quoc-Huy Vu. Semi-quantum copy-protection and more. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part IV*, volume 14372 of *LNCS*, pages 155–182. Springer, Cham, November / December 2023. 2

Gao15. Jingliang Gao. Quantum union bounds for sequential projective measurements. *Phys. Rev. A*, 92:052331, Nov 2015. 8

GGM84. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions (extended abstract). In *25th FOCS*, pages 464–479. IEEE Computer Society Press, October 1984. 25

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. 4, 24

KLYY25. Fuyuki Kitagawa, Jiahui Liu, Shota Yamada, and Takashi Yamakawa. A unified approach to quantum key leasing with a classical lessor. Cryptology ePrint Archive, Report 2025/1871, 2025. 6

KMY25. Fuyuki Kitagawa, Tomoyuki Morimae, and Takashi Yamakawa. A simple framework for secure key leasing. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 217–247. Springer, Cham, May 2025. 2, 3, 4

KN22. Fuyuki Kitagawa and Ryo Nishimaki. Functional encryption with secure key leasing. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 569–598. Springer, Cham, December 2022. 1

KNP25a. Fuyuki Kitagawa, Ryo Nishimaki, and Nikhil Pappu. Collusion-resistant quantum secure key leasing beyond decryption. Cryptology ePrint Archive, Report 2025/1842, 2025. 6

KNP25b. Fuyuki Kitagawa, Ryo Nishimaki, and Nikhil Pappu. PKE and ABE with collusion-resistant secure key leasing. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 35–68. Springer, Cham, August 2025. 3, 6

LG22. Zhenning Liu and Alexandru Gheorghiu. Depth-efficient proofs of quantumness. *Quantum*, 6:807, September 2022. 7

LMZ23. Jiahui Liu, Hart Montgomery, and Mark Zhandry. Another round of breaking and making quantum money: How to not build it from lattices, and more. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 611–638. Springer, Cham, April 2023. 4, 27

LZ19. Qipeng Liu and Mark Zhandry. Revisiting post-quantum Fiat-Shamir. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 326–355. Springer, Cham, August 2019. 4, 27

Mah18. Urmila Mahadev. Classical verification of quantum computations. In Mikkel Thorup, editor, *59th FOCS*, pages 259–267. IEEE Computer Society Press, October 2018. 6

MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. 24

NC11. Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, USA, 10th edition, 2011. 8

Pei10. Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 80–97. Springer, Berlin, Heidelberg, August 2010. 24

Por23. Alexander Poremba. Quantum proofs of deletion for learning with errors. In Yael Tauman Kalai, editor, *ITCS 2023*, volume 251, pages 90:1–90:14. LIPIcs, January 2023. 4, 9, 27

PWYZ24. Duong Hieu Phan, Weiqiang Wen, Xingyu Yan, and Jinwei Zheng. Adaptive hardcore bit and quantum key leasing over classical channel from LWE with polynomial modulus. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part IX*, volume 15492 of *LNCS*, pages 185–214. Springer, Singapore, December 2024. 2, 3, 4, 6

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. 6, 24

Sha84. Adi Shamir. Identity-based cryptosystems and signature schemes. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 47–53. Springer, Berlin, Heidelberg, August 1984. 2

Shm22. Omri Shmueli. Public-key quantum money with a classical bank. In Stefano Leonardi and Anupam Gupta, editors, *54th ACM STOC*, pages 790–803. ACM Press, June 2022. 2

SSTX09. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 617–635. Springer, Berlin, Heidelberg, December 2009. 24

Win99. Andreas Winter. Coding theorem and strong converse for quantum channels. *IEEE Transactions on Information Theory*, 45(7):2481–2485, 1999. 13

WZ82. W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982. 1

Zha20. Mark Zhandry. Schrödinger's pirate: How to trace a quantum decoder. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 61–91. Springer, Cham, November 2020. 25