Quantum Proofs of Deletion for Learning with Errors

Alexander Poremba*

California Institute of Technology

August 23, 2022

Abstract

Quantum information has the property that measurement is an inherently destructive process. This feature is most apparent in the principle of complementarity, which states that mutually incompatible observables cannot be measured at the same time. Recent work by Broadbent and Islam (TCC 2020) builds on this aspect of quantum mechanics to realize a cryptographic notion called *certified deletion*. While this remarkable notion enables a classical verifier to be convinced that a (private-key) quantum ciphertext has been deleted by an untrusted party, it offers no additional layer of functionality.

In this work, we augment the proof-of-deletion paradigm with fully homomorphic encryption (FHE). We construct the first fully homomorphic encryption scheme with certified deletion – an interactive protocol which enables an untrusted quantum server to compute on encrypted data and, if requested, to simultaneously prove data deletion to a client. Our scheme has the desirable property that verification of a deletion certificate is *public*; meaning anyone can verify that deletion has taken place. Our main technical ingredient is an interactive protocol by which a quantum prover can convince a classical verifier that a sample from the Learning with Errors (LWE) distribution in the form of a quantum state was deleted. As an application of our protocol, we construct a *Dual-Regev* public-key encryption scheme with certified deletion, which we then extend towards a (leveled) FHE scheme of the same type. We introduce the notion of *Gaussian-collapsing* hash functions – a special case of collapsing hash functions defined by Unruh (Eurocrypt 2016) – and we prove the security of our schemes under the assumption that the Ajtaj hash function satisfies a certain *strong* Gaussian-collapsing property in the presence of leakage.

Our results enable a form of everlasting cryptography and give rise to new privacy-preserving quantum cloud applications, such as private machine learning on encrypted data with certified data deletion.

*aporemba@caltech.edu

Contents

1 Introduction

Data protection has become a major challenge in the age of cloud computing and artificial intelligence. The European Union, Argentina, and California recently introduced new data privacy regulations which grant individuals the right to request the deletion of their personal data by *media companies* and other *data collectors* – a legal concept that is commonly referred to as the *right to be forgotten* [GGV20]. While new data privacy regulations have been put into practice in several jurisdictions, formalizing data deletion remains a fundamental challenge for cryptography. A key question, in particular, prevails:

How can we certify that user data stored on a remote cloud server has been deleted?

Without any further assumptions, the task is clearly impossible to realize in conventional cloud computing. This is due to the fact that there is no way of preventing the data collector from generating and distributing additional copies of the user data. Although it impossible to achieve in general, *proofs-of-secure-erasure* [PT10, DKW11] can achieve a limited notion of data deletion under *bounded memory assumptions*. Recently, Garg, Goldwasser and Vasudevan [GGV20] proposed rigorous definitions that attempt to formalize the *right to be forgotten* from the perspective of classical cryptography. However, a fundamental challenge in the work of Garg et al. [GGV20] lies in the fact that the data collector is always assumed to be *honest*, which clearly limits the scope of the formalism.

A recent exciting idea is to use quantum information in the context of data privacy [CRW19, BI20]. Contrary to classical data, it is fundamentally impossible to create copies of an unknown quantum state thanks to the *quantum no-cloning theorem* [WZ82]. Building on the work of Coiteux-Roy and Wolf [CRW19], Broadbent and Islam [BI20] proposed a quantum encryption scheme which enables a user to certify the deletion of a quantum ciphertext. Unlike classical proofs-of-secure-erasure, this notion of certified deletion is achievable unconditionally in a fully malicious adversarial setting [BI20]. All prior protocols for certified deletion enable a client to delegate data in the form of plaintexts and ciphertexts with no additional layer of functionality. A key question raised by Broadbent and Islam [BI20] is the following:

Can we enable a remote cloud server to compute on encrypted data, while simultaneously allowing the server to prove data deletion to a client?

This cryptographic notion can be seen as an extension of fully homomorphic encryption schemes [RAD78, Gen09, BV11] which allow for arbitrary computations over encrypted data. Prior work on certified deletion makes use of very specific encryption schemes that seem incompatible with such a functionality; for example, the private-key encryption scheme of Broadbent and Islam [BI20] requires a classical *one-time pad*, whereas the authors in [HMNY21b] use a particular *hybrid encryption* scheme in the context of public-key cryptography. While homomorphic encryption enables a wide range of applications including private queries to a search engine and machine learning classification on encrypted data [BPTG14], a fundamental limitation remains: once the protocol is complete, the cloud server is still in possession of the client's encrypted data. This may allow adversaries to break the encryption scheme retrospectively, i.e. long after the execution of the protocol. This potential threat especially concerns data which is required to remain confidential for many years, such as medical records or government secrets.

Fully homomorphic encryption with certified deletion seeks to address this limitation as it allows a quantum cloud server to compute on encrypted data while simultaneously enabling the server to prove data deletion to a client, thus effectively achieving a form of *everlasting security* [MQU07, HMNY21a].

1.1 Main results

Our contributions are the following.

Quantum superpositions of LWE samples. We use Gaussian states to encode samples from the Learning with Errors (LWE) distribution [Reg05] for the purpose of *certified deletion* while simultaneously preserving their full cryptographic functionality. Because verification of a deletion certificate amounts to checking whether it is a solution to the *(inhomogenous) short integer solution* problem [Ajt96], our encoding results in encryption schemes with certified deletion which are publicly verifiable – in contrast to prior work based on hybrid encryption and BB84 states [BI20, HMNY21a]. Our technique suggests a generic template for *certified deletion* protocols which can be applied to many other cryptographic primitives based on LWE.

Gaussian-collapsing hash functions. To analyze the security of our quantum encryption schemes based on Gaussian states, we introduce the notion of *Gaussian-collapsing* hash functions – a special class of so-called *collapsing* hash functions defined by Unruh [Unr15]. Informally, a hash function h is *Gaussian-collapsing* if it is computationally difficult to distinguish a superposition of Gaussian-weighted pre-images under h from a single (measured) pre-image. We prove that the *Ajtaj collision-resistant hash function* [Ajt96] is Gaussian-collapsing assuming the quantum subexponential hardness of decisional LWE.

Dual-Regev public-key encryption with certified deletion. Using Gaussian superpositions, we construct a public-key encryption scheme with certified deletion which is based on the *Dual-Regev* scheme introduced by Gentry, Peikert and Vaikuntanathan [GPV07]. We prove the security of our scheme under the assumption that Ajtaj's hash function satisfies a certain strong Gaussian-collapsing property in the presence of leakage.

(Leveled) fully homomorphic encryption with certified deletion. We construct the first (leveled) fully homomorphic encryption (FHE) scheme with certified deletion based on our aforementioned *Dual-Regev* encryption scheme with the identical security guarantees. Our FHE scheme is based on the (classical) *dual homomorphic encryption* scheme used by Mahadev [Mah18], which is a variant of the FHE scheme by Gentry, Sahai and Waters [GSW13]. Our protocol supports the evaluation of polynomial-sized Boolean circuits on encrypted data and, if requested, also enables the server to prove data deletion to a client.

1.2 Overview

How can we certify that sensitive information stored on a remote cloud server has been deleted? Remarkably, quantum information allows us to achieve the notion of *certified deletion* using the principle of complementarity; in other words, by encoding information in two mutually incompatible bases.

Broadbent and Islam [BI20] construct a private-key quantum encryption scheme with a rigorous notion of certified deletion using a BB84-type protocol that closely resembles the standard quantum key distribution protocol [BB84, TL17]. There, the ciphertext (without the optional quantum error correction part) consists of random BB84 states $|x^\theta\rangle = H^{\theta_1}|x_1\rangle \otimes \cdots \otimes H^{\theta_n}|x_n\rangle$ together with a one-time pad encryption of the form $f(x_{|\theta_i=0}) \oplus m \oplus u$, where u is a random string (i.e. a one-time pad key), f is a two-universal hash function and $x_{|\theta_i=0}$ is the substring of x to which no Hadamard gate is applied. The main idea behind the scheme is that the information which is necessary to decrypt is encoded in the *computational basis*, whereas deletion requires a *Hadamard basis* measurement. Therefore, if the verification of a deletion certificate is successful, $x_{|\theta_i=0}$ must have high entropy, and thus $f(x_{|\theta_i=0})$ is statistically close to uniform (i.e. f serves

as an extractor). The scheme in [BI20] achieves the notion of *certified deletion security*: once the ciphertext is successfully deleted, the plaintext m remains hidden even if the private key (θ, f, u) is later revealed.

Using a standard *hybrid encryption scheme*, Hiroka, Morimae, Nishimaki and Yamakawa [HMNY21b] extended the scheme in [BI20] to both public-key and attribute-based encryption with certified deletion via the notion of *receiver non-committing* (RNC) encryption [JL00, CFGN96]. For example, in order to obtain a public-key encryption scheme with certified deletion, one simply outputs a ciphertext of the [BI20] scheme together with a classical (non-committing) public-key encryption of its private key, i.e.

$$\mathsf{CT} \leftarrow \left(|x^\theta\rangle, \quad f(x_{|\theta_i=0}) \oplus m \oplus u, \quad \mathsf{RNC.Enc}_{\mathsf{pk}}(\theta||f||u) \right).$$

Given access to the RNC secret key sk, it is therefore possible to decrypt CT in order to obtain the plaintext m. Crucially, the hybrid encryption scheme also inherits the certified deletion property of the [BI20] scheme; namely, once deletion has taken place, m remains hidden even if the RNC secret key sk is later revealed. The security proof in [HMNY21b] relies heavily on the fact that the classical public-key encryption is *non-committing*, i.e. it comes with the ability to equivocate ciphertexts to encryptions of arbitrary plaintexts. As a complementary result, the authors also gave a public-key encryption scheme with certified deletion which is *publicly verifiable* assuming the existence of one-shot signatures and extractable witness encryption. This property enables anyone to verify a deletion certificate using a publicly available verification key.

All prior protocols for certified deletion enable a client to delegate data in the form of plaintexts and ciphertexts with no additional layer of functionality. In this work, we answer a question raised by Broadbent and Islam [BI20] affirmatively, namely whether it is possible to construct a *homomorphic* quantum encryption scheme with certified deletion. This cryptographic notion is remarkably powerful as it would allow a quantum cloud server to compute on encrypted data, while simultaneously enabling the server to prove data deletion to a client. So far, however, none of the encryption schemes with certified deletion can enable such a functionality. Worse yet, the hybrid encryption paradigm appears fundamentally insufficient in order to construct homomorphic encryption with certified deletion: once we instantiate the (non-committing) public-key encryption scheme with a (classical) fully homomorphic encryption (FHE) encryption scheme instead, anyone can simply run the following homomorphic evaluation procedure (in superposition) to compute

$$\left(|x^\theta\rangle, \quad f(x_{|\theta_i=0}) \oplus m \oplus u, \quad \mathsf{FHE.Enc}_{\mathsf{pk}}(\theta||f||u) \right) \quad \overset{\mathsf{Eval}}{\mapsto} \quad \mathsf{FHE.Enc}_{\mathsf{pk}}(m).$$

Assuming that the FHE scheme is *correct*, this step can be performed without disturbing the BB84 state $|x^\theta\rangle$ in the process.[1] Notice that the classical ciphertext $\mathsf{FHE.Enc}_{\mathsf{pk}}(m)$ is now completely decoupled from everything else. In particular, since the BB84 state $|x^\theta\rangle$ remains intact, it is possible to prove deletion and to simultaneously recover m once the secret key is revealed. Because homomorphic encryption schemes are *malleable* by design, it seems fundamentally impossible for an encryption scheme to be homomorphic and non-committing at the same time. Therefore, to construct a *truly* homomorphic encryption scheme with certified deletion, an entirely new approach is necessary. Our techniques deviate from the hybrid encryption paradigm of previous works [BI20, HMNY21a] and allow us to construct the *first* homomorphic quantum encryption scheme with certified deletion which has the desirable feature of being publicly verifiable.

Quantum superpositions of LWE samples. The *Learning with Errors* (LWE) problem was introduced by Regev [Reg05] and serves as the primary basis of hardness for post-quantum cryptosystems, mainly due to its tight connection with worst-case approximation problems over Euclidean lattices.

[1] For example, by relying on the so-called *Almost As Good As New Lemma* [Aar16].

The problem is described as follows. Let $n, m \in \mathbb{N}$ and $q \geq 2$ be a prime modulus, and $\alpha \in (0,1)$ be a noise ratio parameter. In its decisional formulation, the $\mathsf{LWE}^m_{n,q,\alpha q}$ problem asks to distinguish between a sample $(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}, \mathbf{s}\mathbf{A} + \mathbf{e} \pmod q)$ from the LWE distribution and a uniformly random sample $(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}, \mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m)$. Here, $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ is chosen uniformly random and $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ is an error vector which is sampled according to the discrete Gaussian distribution $D_{\mathbb{Z}^m,\alpha q}$. The latter distribution assigns probability proportional to $\varrho_r(\mathbf{x}) = e^{-\pi\|\mathbf{x}\|^2/r^2}$ to every lattice point $\mathbf{x} \in \mathbb{Z}^m$, for $r = \alpha q > 0$.

How can we certify that a (possibly malicious) prover has deleted a sample from the LWE distribution? The main technical insight of our work is that one can encode LWE samples as *quantum superpositions* for the purpose of certified deletion while simultaneously preserving their full cryptographic functionality. Superpositions of LWE samples have been considered by Grilo, Kerenidis and Zijlstra [GKZ19] in the context of quantum learning theory and by Alagic, Jeffery, Ozols and Poremba [AJOP20], as well as by Chen, Liu and Zhandry [CLZ21], in the context of quantum cryptanalysis of LWE-based cryptosystems.

Let us now describe the main idea behind our constructions. Consider the Gaussian superposition,[2]

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

Here, we let $\sigma = 1/\alpha$ and we use $\mathbb{Z}_q^m$ to represent $\mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$. By measuring system Y in the computational basis with outcome $\mathbf{y} \in \mathbb{Z}_q^n$, the state $|\hat{\psi}\rangle$ *collapses* into the quantum superposition

$$|\hat{\psi}_\mathbf{y}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle . \tag{1}$$

Note that the state $|\hat{\psi}_\mathbf{y}\rangle$ is now a superposition of *short* Gaussian-weighted solutions $\mathbf{x} \in \mathbb{Z}_q^m$ subject to the constraint $\mathbf{A}\cdot\mathbf{x} = \mathbf{y} \pmod q$. In other words, by measuring the above state in the computational basis, we obtain a solution to the so-called *(inhomogenous) short integer solution* (ISIS) problem specified by $(\mathbf{A}, \mathbf{y})$ (see Definition 13). The quantum state $|\hat{\psi}_\mathbf{y}\rangle$ in Eq. (1) has a surprising *duality property*; namely, by applying the (inverse) q-ary quantum Fourier transform we obtain the state

$$|\psi_\mathbf{y}\rangle = \sum_{\mathbf{s}\in\mathbb{Z}_q^n} \sum_{\mathbf{e}\in\mathbb{Z}_q^m} \varrho_{\frac{q}{\sigma}}(\mathbf{e}) \, \omega_q^{-\langle \mathbf{s},\mathbf{y}\rangle} \, |\mathbf{s}\mathbf{A} + \mathbf{e} \pmod q\rangle , \tag{2}$$

where $\omega_q = e^{2\pi i/q}$ is the primitive q-th root of unity. We make this statement more precise in Lemma 16. Throughout this work, we will refer to $|\psi_\mathbf{y}\rangle$ and $|\hat{\psi}_\mathbf{y}\rangle$ as the *primal* and *dual* Gaussian state, respectively. Notice that the resulting state $|\psi_\mathbf{y}\rangle$ is now a quantum superposition of samples from the LWE distribution. This relationship was first observed in the work of Stehlé et al. [SSTX09] who gave quantum reduction from SIS to LWE based on Regev's reduction [Reg05], and was later implicitly used by Roberts [Rob19] and Kitagawa et al. [KNY21] to construct quantum money and secure software leasing schemes.

Our quantum encryption schemes with certified deletion exploit the fact measurement of $|\psi_\mathbf{y}\rangle$ in the *Fourier basis* yields a short solution to the ISIS problem specified by $(\mathbf{A}, \mathbf{y})$, whereas information which is necessary to decrypt is encoded using LWE samples in the (incompatible) *computational basis*.

[2] A standard tail bound shows that the discrete Gaussian $D_{\mathbb{Z}^m,\sigma}$ is essentially only supported on $\{\mathbf{x} \in \mathbb{Z}^m : \|\mathbf{x}\|_\infty \leq \sigma\sqrt{m}\}$. We choose $\sigma \ll q/\sqrt{m}$ and consider the domain $\mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$ instead. For simplicity, we also ignore that $|\hat{\psi}\rangle$ is not normalized.

Dual-Regev public-key encryption with certified deletion. The key ingredient of our homomorphic encryption scheme with certified deletion is the *Dual-Regev* public-key encryption (PKE) scheme introduced by Gentry, Peikert and Vaikuntanathan [GPV07]. Unlike Regev's original PKE scheme in [Reg05], the Dual-Regev PKE scheme has the property that the ciphertext takes the form of a regular sample from the LWE distribution together with an additive shift $b \cdot \lfloor \frac{q}{2} \rfloor$ that depends on the plaintext $b \in \{0,1\}$. Using Gaussian states, we can encode Dual-Regev ciphertexts for the purpose of certified deletion while simultaneously preserving their cryptographic functionality. The scheme consists of the following efficient algorithms:

- To generate a pair of keys $(\mathsf{sk}, \mathsf{pk})$, sample a random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times (m+1)}$ together with a short trapdoor vector $\mathbf{t} = (\bar{\mathbf{x}}, -1) \in \mathbb{Z}^{m+1}$ such that $\mathbf{A} \cdot \mathbf{t} = \mathbf{0} \pmod q$, and let $\mathsf{pk} = \mathbf{A}$ and $\mathsf{sk} = \mathbf{t}$.
- To encrypt $b \in \{0,1\}$ using the public key $\mathsf{pk} = \mathbf{A}$, choose a random $\mathbf{y} \in \mathbb{Z}_q^n$ and output the pair

$$\mathsf{vk} \leftarrow (\mathbf{A}, \mathbf{y}), \qquad |\mathsf{CT}\rangle \leftarrow \sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \varrho_{\frac{q}{\sigma}}(\mathbf{e})\, \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle}\, |\mathbf{s}\mathbf{A} + \mathbf{e} + (0, \ldots, 0, b \cdot \lfloor \frac{q}{2} \rfloor)\rangle\,,$$

  where vk is a public verification key and $|\mathsf{CT}\rangle$ is the quantum ciphertext.
- To decrypt a ciphertext $|\mathsf{CT}\rangle$ using the secret key sk, measure in the computational basis to obtain an outcome $\mathbf{c} \in \mathbb{Z}_q^{m+1}$, and output 0, if $\mathbf{c}^T \cdot \mathsf{sk} \in \mathbb{Z}_q$ is closer to 0 than to $\lfloor \frac{q}{2} \rfloor$, and output 1, otherwise.

To delete the ciphertext $|\mathsf{CT}\rangle$, we simply perform measurement in the Fourier basis. In Corollary 1, we show that the Fourier transform of the ciphertext $|\mathsf{CT}\rangle$ results in the quantum state

$$|\widehat{\mathsf{CT}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^{m+1}: \\ \mathbf{A}\mathbf{x} = \mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x})\, \omega_q^{\langle \mathbf{x}, (0,\ldots,0,b \cdot \lfloor \frac{q}{2} \rfloor) \rangle}\, |\mathbf{x}\rangle\,. \tag{3}$$

Notice that a Fourier basis measurement of $|\mathsf{CT}\rangle$ necessarily erases all information about the plaintext $b \in \{0,1\}$ and results in a *short* vector $\pi \in \mathbb{Z}_q^{m+1}$ such that $\mathbf{A} \cdot \pi = \mathbf{y} \pmod q$. In other words, to verify a deletion certificate we can simply check whether it is a solution to the ISIS problem specified by the verification key $\mathsf{vk} = (\mathbf{A}, \mathbf{y})$. Our scheme has the desirable property that verification of a certificate π is public; meaning anyone in possession of $(\mathbf{A}, \mathbf{y})$ can verify that $|\mathsf{CT}\rangle$ has been successfully deleted. Moreover, due to the tight connection between worst-case lattice problems and the average-case ISIS problem [MR07, GPV07], it is computationally difficult to produce a valid deletion certificate from $(\mathbf{A}, \mathbf{y})$ alone.

To formalize security, we use the notion of *certified deletion security* (i.e. IND-CPA-CD security) [BI20, HMNY21a] which roughly states that, once deletion of the ciphertext is successful, the plaintext remains hidden even if the secret key is later revealed (see Definition 23).

Unfortunately, proving security from standard assumptions, such as LWE (or ISIS), is highly non-trivial. The problem emerges when we attempt to reduce the IND-CPA-CD security of our Dual-Regev public-key encryption scheme with certified deletion to the LWE (or ISIS) problem. In order to simulate the IND-CPA-CD game successfully, we have to eventually forward a short trapdoor vector $\mathbf{t} \in \mathbb{Z}^{m+1}$ (i.e. the secret key) to the adversary once deletion has taken place. Notice, however, that the reduction has no way of obtaining a short trapdoor vector $\mathbf{t}$ such that $\mathbf{A} \cdot \mathbf{t} = \mathbf{0} \pmod q$ as it is trying to break the underlying LWE (or ISIS) problem with respect to $\mathbf{A}$ in the first place (!) Recently, Hiroka, Morimae, Nishimaki and Yamakawa [HMNY21a] managed to overcome similar technical difficulties using the notion of *receiver non-committing* (RNC) encryption [JL00, CFGN96] in the context of *hybrid encryption* in order to produce

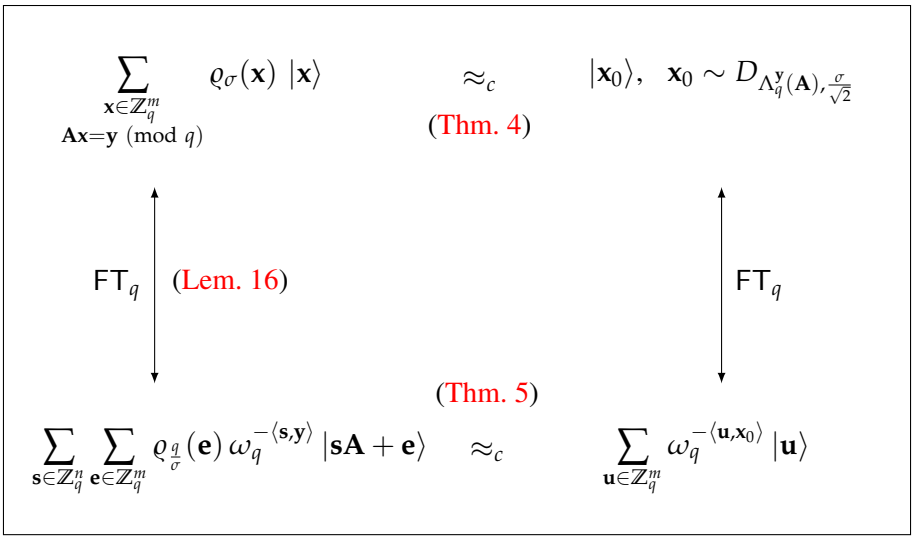

Figure 1: Technical overview of the main quantum states and their properties used throughout this work. The computational indistinguishability property holds under the (subexponential) quantum hardness of the (decisional) LWE assumption (Definition 15). Here, $\Lambda_q^{\mathbf{y}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\cdot\mathbf{x} = \mathbf{y} \pmod q\}$ denotes a particular coset of the q-ary lattice $\Lambda_q^{\perp}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m :\ \mathbf{A}\cdot\mathbf{x} = \mathbf{0} \pmod q\}$ defined in Section 2.5.

a *fake* secret key. In our case, we cannot rely on similar techniques involving RNC encryption as it seems difficult to reconcile with homomorphic encryption, which is the main focus of this work.

To prove the security of our scheme, we instead rely on a new conjecture which states that Ajtaj's hash function $h_{\mathbf{A}}(\mathbf{x}) = \mathbf{A}\cdot\mathbf{x} \pmod q$ satisfies a certain strong *collapsing property* in the presence of leakage.

Gaussian-collapsing hash functions. Unruh [Unr15] introduced the notion of collapsing hash functions in his seminal work on computationally binding quantum commitments. Informally, a hash function h is called *collapsing* if it is computationally difficult to distinguish between a superposition of pre-images, i.e. $\sum_{\mathbf{x}:\,h(\mathbf{x})=\mathbf{y}} \alpha_{\mathbf{x}}\,|\mathbf{x}\rangle$, and a single measured pre-image $|\mathbf{x}_0\rangle$ such that $h(\mathbf{x}_0) = \mathbf{y}$. Motivated by the properties of the dual Gaussian state in Eq. (1), we consider a special class of hash functions which are *collapsing* with respect to Gaussian superpositions. We say that a hash function h is σ-*Gaussian-collapsing* (formally defined in Definition 19), for some $\sigma > 0$, if the following states are computationally indistinguishable:

$$\sum_{\mathbf{x}:\,h(\mathbf{x})=\mathbf{y}} \varrho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle \quad \approx_c \quad |\mathbf{x}_0\rangle\,,\ \text{s.t.}\ \ h(\mathbf{x}_0) = \mathbf{y}.$$

Here, $\mathbf{x}_0$ is the result of a computational basis measurement of the the Gaussian superposition (on the left). Notice that any collapsing hash function h is necessarily also *Gaussian-collapsing*, since a superposition of Gaussian-weighted vectors constitutes a special class of inputs to h. Liu and Zhandry [LZ19] implicitly showed that the *Ajtaj hash function* $h_{\mathbf{A}}(\mathbf{x}) = \mathbf{A}\cdot\mathbf{x} \pmod q$ is collapsing – and thus *Gaussian-collapsing* – via the notion of *lossy functions* by assuming the superpolynomial hardness of (decisional) LWE. In Theorem 4, we give a simple and direct proof that the Ajtaj hash function is Gaussian-collapsing assuming (decisional) LWE, which might be of independent interest.

The fact Ajtaj's hash function is Gaussian-collapsing has several implications for the security of our Dual-Regev public-key encryption scheme with certified deletion. Because our Dual-Regev ciphertext $|\mathsf{CT}\rangle$

corresponds to the Fourier transform of the dual Gaussian state in Eq. (3), the Gaussian-collapsing property immediately implies the semantic (i.e., IND-CPA) security under the hardness of decisional LWE (see Theorem 5). We refer to Figure 1 for an overview of the primal and dual Gaussian states and their properties.

To prove the stronger notion of IND-CPA-CD security of our Dual-Regev scheme with certified deletion, we have to show that, once deletion has taken place, the plaintext remains hidden even if the secret key (i.e., a short trapdoor vector $\mathbf{t}$ in the kernel of $\mathbf{A}$) is later revealed. In other words, it is sufficient to show that Ajtaj's hash function satisfies a particular *strong* Gaussian-collapsing property in the presence of leakage; namely, once a computationally bounded adversary $\mathcal{A}$ produces a valid short certificate π with the property that $\mathbf{A} \cdot \pi = \mathbf{y} \pmod q$, then $\mathcal{A}$ cannot tell whether the input at the beginning of the experiment corresponded to a Gaussian superposition of pre-images or a single (measured) pre-image, even if $\mathcal{A}$ later receives a short trapdoor vector $\mathbf{t}$ in the kernel of $\mathbf{A}$. Here, it is crucial that $\mathcal{A}$ receives the trapdoor vector $\mathbf{t}$ only *after* $\mathcal{A}$ provides a valid pre-image witness π, otherwise $\mathcal{A}$ could trivially distinguish the two states by applying the Fourier transform and using the trapdoor $\mathbf{t}$ to check whether the outcome corresponds to a superposition of LWE samples (rather than a uniform superposition). Unfortunately, we currently do not know how to prove the *strong* Gaussian-collapsing property of the Ajtaj hash function from standard assumptions (such as LWE or ISIS); instead, we choose to formalize it as a simple and falsifiable conjecture in Conjecture 5.2.

To see why Conjecture 5.2 is plausible, consider the following natural attack. Given as input either a Gaussian superposition of pre-images or a single (measured) pre-image, we perform the quantum Fourier transform, reversibly shift the outcome by a fresh LWE sample[3] and store the result in an auxiliary register. If the input corresponds to a superposition, we obtain a separate LWE sample which is *re-randomized*, whereas if the input is a single (measured) pre-image, the outcome remains random. Hence, if the aforementioned procedure were to succeed without disturbing the initial quantum state, we could potentially provide a valid certificate π and also distinguish the auxiliary system with access to the trapdoor. However, by shifting the initial state by another LWE sample, we have necessarily entangled the two systems in a way that prevents us from obtaining a valid certificate via the required Fourier basis measurement. We make this fact more precise in Section 4, where we prove a general *uncertainty relation for Fourier basis projections* (Theorem 3) that rules out a large class of attacks, including the *shift-by-*LWE*-sample* attack described above.

We prove the following result in Theorem 7 (assuming that Conjecture 5.2 holds):

Theorem (informal): *The Dual-Regev* PKE *scheme with certified deletion (see Construction 1) is* IND-CPA-CD*-secure under the strong Gaussian-collapsing assumption in Conjecture 5.2.*

Next, we extend our Dual-Regev public-key encryption scheme with certified deletion towards a (leveled) FHE scheme of the same type.

Dual-Regev fully homomorphic encryption with certified deletion. Our (leveled) FHE scheme with certified deletion is based on the (classical) Dual-Regev leveled FHE scheme used by Mahadev [Mah18] – a variant of the scheme due to Gentry, Sahai and Waters [GSW13]. Let $n, m \in \mathbb{N}$, let $q \geq 2$ be a prime modulus, and let $\alpha \in (0,1)$ be the noise ratio with $\sigma = 1/\alpha$. Let $N = (n+1)\lceil \log q \rceil$ and let $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ denote the *gadget matrix* (defined in Section 9.1) designed to convert a binary representation of a vector back to its $\mathbb{Z}_q$ representation. The scheme consists of the following efficient algorithms:

- To generate a pair of keys $(\mathsf{sk}, \mathsf{pk})$, sample a random matrix $\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}$ together with a short

[3] To *smudge* the Gaussian error of the initial superposition, we can choose an error from a discrete Gaussian distribution which has a significantly larger standard deviation.

trapdoor vector $\mathbf{t} = (\bar{\mathbf{x}}, -1) \in \mathbb{Z}^{m+1}$ such that $\mathbf{t} \cdot \mathbf{A} = \mathbf{0} \pmod{q}$, and let $\mathsf{pk} = \mathbf{A}$ and $\mathsf{sk} = \mathbf{t}$.

- To encrypt a bit $x \in \{0,1\}$ using the public key $\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}$, choose random matrix $\mathbf{Y} \in \mathbb{Z}_q^{n\times N}$ composed of rows $\mathbf{y}_1, \ldots, \mathbf{y}_N \in \mathbb{Z}_q^n$ and output a public verification key and ciphertext pair

$$\mathsf{vk} \leftarrow (\mathbf{A}, \mathbf{Y}), \quad |\mathsf{CT}\rangle \leftarrow \sum_{\mathbf{S}\in\mathbb{Z}_q^{n\times N}} \sum_{\mathbf{E}\in\mathbb{Z}_q^{(m+1)\times N}} \varrho_{q/\sigma}(\mathbf{E})\, \omega_q^{-\mathrm{Tr}[\mathbf{S}^T\mathbf{Y}]}\, |\mathbf{A}\cdot\mathbf{S} + \mathbf{E} + x\cdot\mathbf{G}\rangle\,,$$

  where $(\mathbf{g}_1, \ldots, \mathbf{g}_N)$ are the rows of the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$.

- To decrypt a quantum ciphertext $|\mathsf{CT}\rangle$ using the secret key sk, measure in the computational basis to obtain an outcome $\mathbf{C} \in \mathbb{Z}_q^{(m+1)\times N}$ and compute $c = \mathsf{sk}^T \cdot \mathbf{c}_N \in \mathbb{Z}_q$, where $\mathbf{c}_N \in \mathbb{Z}_q^{m+1}$ is the N-th column of $\mathbf{C}$, and then output 0, if c is closer to 0 than to $\lfloor \frac{q}{2} \rfloor$, and output 1, otherwise.

We remark that deletion and verification take place as in our Dual-Regev scheme with certified deletion.

Let us now describe how to perform homomorphic operations on the encrypted data. Our FHE scheme supports the evaluation of polynomial-sized Boolean circuits consisting entirely of NOT-AND (NAND) gates, which are universal for classical computation. Recall that the (classical) Dual-Regev FHE scheme supports the homomorphic evaluation of a NAND gate in the following sense. If CT_0 and CT_1 are ciphertexts that encrypt two bits x_0 and x_1, respectively, then the outcome $\mathsf{CT} = \mathbf{G} - \mathsf{CT}_0 \cdot \mathbf{G}^{-1}(\mathsf{CT}_1) \pmod{q}$ is an encryption of $\mathsf{NAND}(x_0, x_1) = 1 - x_0 \cdot x_1$. Moreover, the new ciphertext CT maintains the form of an LWE sample with respect to the same public key pk, albeit for a new LWE secret and a new (non-necessarily Gaussian) noise term of bounded magnitude. This property is crucial, as knowledge of the secret key sk still allows for the decryption of the ciphertext CT once a NAND gate has been applied (see Section 9.1).

Inspired by the classical homomorphic NAND operation, we define an analogous quantum operation U_{NAND} in Definition 28 which allows us to apply a NAND gate directly onto Gaussian states. Consider two ciphertexts $|\mathsf{CT}_0\rangle$ and $|\mathsf{CT}_1\rangle$ in systems C_0 and C_1, respectively. Applying the homomorphic NAND gate via the unitary U_{NAND} results in an output state CT in systems $C_0C_1C_{\mathsf{out}}$ such that

$$U_{\mathsf{NAND}}: \quad |\mathsf{CT}_0\rangle_{C_0} \otimes |\mathsf{CT}_1\rangle_{C_1} \otimes |\mathbf{0}\rangle_{C_{\mathsf{out}}} \quad \mapsto \quad |\mathsf{CT}\rangle_{C_0C_1C_{\mathsf{out}}}\,. \tag{4}$$

Just as in the (classical) Dual-Regev FHE scheme, the basis states of the state $|\mathsf{CT}\rangle$ in system C_{out} maintain the form of an LWE sample with a new bounded noise vector. Therefore, in principle, it should be possible to measure the outcome in system C_{out} in order to learn the ciphertext that corresponds to an encryption of $\mathsf{NAND}(x_0, x_1) = 1 - x_0 \cdot x_1$. Notice, however, that the new ciphertext $|\mathsf{CT}\rangle$ is now a highly entangled state since the unitary operation U_{NAND} induces entanglement between the LWE secrets and Gaussian error terms of the superposition. This raises the following question: How can a quantum server perform homomorphic computations and, if requested, to afterwards prove data deletion to a client? In some sense, applying a single homomorphic NAND gates breaks the structure of the Gaussian states in a way that makes it impossible to perform the correct Fourier basis measurement required for a proof of deletion.

Our solution to the problem involves a single additional round of interaction between the quantum server (the prover) and the client (the verifier) in order to prove deletion. After performing the Boolean circuit C via a sequence of U_{NAND} gates starting from the ciphertext $|\mathsf{CT}\rangle = |\mathsf{CT}_1\rangle \otimes \cdots \otimes |\mathsf{CT}_\ell\rangle$ in system C_{in} which corresponds to an encryption of $x = (x_1, \ldots, x_\ell) \in \{0,1\}^\ell$, the prover simply sends the quantum system C_{out} containing an encryption of $C(x)$ to the verifier. Then, using the secret key sk (i.e., a trapdoor for the public matrix pk), it is possible for the verifier to *extract* the outcome $C(x)$ from the system C_{out}

with overwhelming probability without significantly damaging the state. By the *Almost As Good As New Lemma* [Aar16] (see Lemma 1), it is possible to rewind the procedure in a way that results in a state which is negligibly close to the original state in system C_{out}. At this step of the protocol, the verifier has learned the outcome of the homomorphic application of the circuit C while the prover is still in possession of a large number of auxiliary systems (denoted by C_{aux}) which mark intermediate applications of the gate U_{NAND}. In order to allow for a quantum proof of deletion, the verifier must now return the system C_{out} to the prover. Having access to all three systems $C_{\mathsf{in}}C_{\mathsf{aux}}C_{\mathsf{out}}$, the prover is then able to undo the sequence of homomorphic NAND gates in order to return to the original product state in system C_{in} (up to negligible trace distance). Since the ciphertext in the prover's possession is now approximately a simple product of Gaussian states, the prover can perform a Fourier basis measurement of systems C_{in}, as required. Once the protcol is complete, it is therefore possible for the client to know $C(x)$ and to be convinced that data deletion has taken place.

In terms of security, our FHE scheme with certified deletion inherits the same security guarantees as our Dual-Regev PKE scheme with certified deletion. We prove the following in Theorem 10.

Theorem (informal): *Our Dual-Regev (leveled)* FHE *scheme with certified deletion (Construction 3) is* IND-CPA-CD*-secure under the strong Gaussian-collapsing assumption in Conjecture 5.2.*

Open problems. Our results leave open many interesting future research directions. For example, is it possible to prove Conjecture 5.2 – and thus the IND-CPA-CD security of our constructions – from the hardness of LWE or ISIS? Another interesting direction is the following. Since the verification of our proofs of deletion only requires classical computational capabilities, this leaves open the striking possibility that all communication that is required for fully homomorphic encryption with certified deletion can be dequantized entirely, similar to work of Mahadev [Mah18] on delegating quantum computations, as well as recent work on classically-instructed parallel remote state preparation by Gheorghiu, Metger and Poremba [GMP22].

1.3 Applications

Data retention and the right to be forgotten. The European Union, Argentina, and California recently introduced new data privacy regulations – often referred to as the *right to be forgotten* [GGV20] – which grant individuals the right to request the deletion of their personal data by media companies. However, formalizing data deletion still remains a fundamental challenge for cryptography. Our fully homomorphic encryption scheme with certified deletion achieves a rigorous notion of *long-term data privacy*: it enables a remote quantum cloud server to compute on encrypted data and – once it is deleted and publicly verified – the client's data remains safeguarded even in the case of a future leak that reveals the secret key.

Private machine learning on encrypted data. Machine learning algorithms are used for wide-ranging classification tasks, such as medical predictions, spam detection and face recognition. While homomorphic encryption enables a form of privacy-preserving machine learning [BPTG14], a fundamental limitation remains: once the protocol is complete, the cloud server is still in possession of the client's encrypted data. This threat especially concerns data which is required to remain confidential for many years. Our results remedy this situation by enabling private machine learning on encrypted data with certified data deletion.

Everlasting cryptography. Assuming that the server has not broken the computational assumption before data deletion has taken place, our results could potentially transform a long-term LWE assumption [Reg05] into a temporary one, and thus effectively achieve a form of *everlasting security* [MQU07, HMNY21a].

1.4 Related work

The first work to formalize a notion resembling *certified deletion* is due to Unruh [Unr13] who proposed a quantum timed-release encryption scheme that is *revocable*. The protocol allows a user to *return* the ciphertext of a quantum timed-release encryption scheme, thereby losing all access to the data. Unruh's security proof exploits the *monogamy of entanglement* in order to guarantee that the quantum revocation process necessarily erases all information about the plaintext. Subsequently, Coladangelo, Majenz and Poremba [CMP20] adapted this property to *revocable* programs in the context of *secure software leasing*, a weaker notion of *quantum copy-protection* which was proposed by Ananth and La Placa [AP20].

Fu and Miller [FM18] gave the first quantum protocol that proves deletion of a single bit using classical interaction alone. Subsequently, Coiteux-Roy and Wolf [CRW19] proposed a QKD-like conjugate coding protocol that enables certified deletion of a classical plaintext, albeit without a complete security proof. Coiteux-Roy and Wolf also coined the term *privacy delegation* as the means to delegate information to a remote quantum server in a way that prevents the leakage of user data. By design, privacy delegation cannot make eavesdropping impossible – it merely makes it possible for a verifier to be convinced that deletion has taken place. Broadbent and Islam [BI20] construct a quantum encryption scheme with certified deletion whose security proof is similar to that of QKD protocol [BB84, TL17]. The notion of *certified deletion* proposed by Broadbent and Islam is information-theoretic and does not take computational assumptions into account. Subsequently, Hiroka, Morimae, Nishimaki and Yamakawa [HMNY21b] extended the scheme in [BI20] to public-key and attribute-based encryption by using a *hybrid encryption scheme* in combination with *receiver non-committing* (RNC) encryption [JL00, CFGN96]. Hiroka, Morimae, Nishimaki and Yamakawa [HMNY21a] studied *certified everlasting zero-knowledge proofs* for QMA via the notion of *everlasting security* which was first formalized by Müller-Quade and Unruh [MQU07]. A recent paper by Coladangelo, Liu, Liu and Zhandry [CLLZ21] introduces *subspace coset states* in the context of unclonable crytography in a way that loosely resembles our use of primal and dual Gaussian states.

Acknowledgments. The author would like to thank Urmila Mahadev for pointing out an attack on an earlier version of our protocols, and for the idea behind the proof of Theorem 4. The author would also like to thank Thomas Vidick and Vinod Vaikuntanathan for many insightful discussions. The author is also grateful for many useful comments and suggestions made by anonymous reviewers. The author is partially supported by AFOSR YIP award number FA9550-16-1-0495 and the Institute for Quantum Information and Matter (an NSF Physics Frontiers Center; NSF Grant PHY-1733907), and is also grateful for the hospitality of the Simons Institute for the Theory of Computing, where part of this research was carried out.

2 Preliminaries

Notation. For $x \in \mathbb{C}^n$, we denote the ℓ^2 norm by $\|\mathbf{x}\|_2$. For $\mathbf{x} \in \mathbb{Z}^n$, we occasionally also use the max norm $\|\mathbf{x}\|_\infty = \max_i |x_i|$. We denote the expectation value of a random variable X which takes values in $\mathcal{X}$ by $\mathbb{E}[X] = \sum_{x\in\mathcal{X}} x \Pr[X = x]$. The notation $x \xleftarrow{\$} \mathcal{X}$ denotes sampling of x uniformly at random from $\mathcal{X}$, whereas $x \sim D$ denotes sampling of an element x according to the distribution D. We call a non-negative real-valued function $\mu : \mathbb{N} \to \mathbb{R}^+$ negligible if $\mu(n) = o(1/p(n))$, for every polynomial $p(n)$. Given an integer $m \in \mathbb{N}$ and modulus $q \geq 2$, we represent elements in $\mathbb{Z}_q^m$ as integers $\mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$.

2.1 Quantum computation

For a comprehensive overview of quantum computation, we refer to the introductory texts [NC11, Wil13]. We denote a finite-dimensional complex Hilbert space by $\mathcal{H}$, and we use subscripts to distinguish between different systems (or registers). For example, we let $\mathcal{H}_A$ be the Hilbert space corresponding to a system A. The tensor product of two Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$ is another Hilbert space denoted by $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. The Euclidean norm of a vector $|\psi\rangle \in \mathcal{H}$ over the finite-dimensional complex Hilbert space $\mathcal{H}$ is denoted as $\|\psi\| = \sqrt{\langle\psi|\psi\rangle}$. Let $L(\mathcal{H})$ denote the set of linear operators over $\mathcal{H}$. A quantum system over the 2-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^2$ is called a *qubit*. For $n \in \mathbb{N}$, we refer to quantum registers over the Hilbert space $\mathcal{H} = \left(\mathbb{C}^2\right)^{\otimes n}$ as n-qubit states. More generally, we associate *qudits* of dimension $d \geq 2$ with a d-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^d$. We use the word *quantum state* to refer to both pure states (unit vectors $|\psi\rangle \in \mathcal{H}$) and density matrices $\varrho \in \mathcal{D}(\mathcal{H})$, where we use the notation $\mathcal{D}(\mathcal{H})$ to refer to the space of positive semidefinite matrices of unit trace acting on $\mathcal{H}$. For simplicity, we frequently consider *subnormalized states*, i.e. states in the space of positive semidefinite operators over $\mathcal{H}$ with trace norm not exceeding 1, denoted by $\mathcal{S}_{\leq}(\mathcal{H})$. The *trace distance* of two density matrices $\varrho, \sigma \in \mathcal{D}(\mathcal{H})$ is given by

$$\|\varrho - \sigma\|_{\mathrm{tr}} = \frac{1}{2}\mathrm{Tr}\left[\sqrt{(\varrho-\sigma)^\dagger(\varrho-\sigma)}\right].$$

We frequently use the compact notation $\varrho \approx_\varepsilon \sigma$ which means that there exists some $\varepsilon \in [0,1]$ such that $\|\varrho - \sigma\|_{\mathrm{tr}} \leq \varepsilon$. The *purified distance* is defined as $P(\varrho,\sigma) = \sqrt{1 - F(\varrho,\sigma)^2}$, where $F(\varrho,\sigma) = \|\sqrt{\varrho}\sqrt{\sigma}\|_1$ denotes the fidelity. The trace distance and the ℓ^2 distance over $(\mathbb{C}^q)^{\otimes m}$ are related via the inequality,

$$\||\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|\|_{\mathrm{tr}} \leq \| \, |\psi\rangle - |\phi\rangle \, \|_2, \qquad \forall \, |\psi\rangle, |\phi\rangle \in (\mathbb{C}^q)^{\otimes m}.$$

A *classical-quantum* (CQ) state $\varrho \in \mathcal{D}(\mathcal{H}_{XB})$ depends on a classical variable in system X which is correlated with a quantum system B. If the classical system X is distributed according to a probability distribution $P_{\mathcal{X}}$ over the set $\mathcal{X}$, then all possible joint states ϱ_{XB} can be expressed as

$$\varrho_{XB} = \sum_{x\in\mathcal{X}} P_{\mathcal{X}}(x)|x\rangle\langle x|_X \otimes \varrho_B^x.$$

Quantum channels and measurements. A quantum channel $\Phi : L(\mathcal{H}_A) \to L(\mathcal{H}_B)$ is a linear map between linear operators over the Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$. Oftentimes, we use the compact notation $\Phi_{A\to B}$ to denote a quantum channel between $L(\mathcal{H}_A)$ and $L(\mathcal{H}_B)$. We say that a channel Φ is *completely positive* if, for a reference system R of arbitrary size, the induced map $\mathbb{1}_R \otimes \Phi$ is positive, and we call it *trace-preserving* if $\mathrm{Tr}[\Phi(X)] = \mathrm{Tr}[X]$, for all $X \in L(\mathcal{H})$. A quantum channel that is both completely positive and trace-preserving is called a quantum CPTP channel. Let $\mathcal{X}$ be a set. A *generalized measurement* on a system A is a set of linear operators $\{M_A^x\}_{x\in\mathcal{X}}$ such that

$$\sum_{x\in\mathcal{X}} (M_A^x)^\dagger (M_A^x) = \mathbb{1}_A.$$

We can represent a measurement as a CPTP map $\mathcal{M}_{A\to X}$ that maps states on system A to measurement outcomes in a register denoted by X. For example, let $\varrho \in \mathcal{D}(\mathcal{H}_{AB})$ be a bipartite state. Then,

$$\mathcal{M}_{A\to X} : \quad \varrho_{AB} \quad \mapsto \quad \sum_{x\in\mathcal{X}} |x\rangle\langle x|_X \otimes \mathrm{tr}_A\left[M_A^x \varrho_{AB} M_A^{x\,\dagger}\right],$$

yields a normalized classical-quantum state. A positive-operator valued measure (POVM) on a quantum system A is a set of Hermitian positive semidefinite operators $\{M_A^x\}_{x\in\mathcal{X}}$ such that

$$\sum_{x\in\mathcal{X}} M_A^x = \mathbb{1}_A.$$

Oftentimes, we identify a POVM $\{M_A^x\}_{x\in\mathcal{X}}$ with an associated generalized measurement $\{\sqrt{M_A^x}\}_{x\in\mathcal{X}}$. The *overlap* c of two POVMs $\{M_A^x\}_{x\in\mathcal{X}}$ and $\{N_A^y\}_{y\in\mathcal{X}}$ acting on a quantum system A is defined by

$$\mathsf{c} = \max_{x,y} \left\| \sqrt{M_A^x}\sqrt{N_A^y} \right\|_\infty^2.$$

We say that two measurements are *mutually unbiased*, if the overlap satisfies $\mathsf{c} = 1/d$, where $d = \dim(\mathcal{H}_A)$ is the dimension of the associated Hilbert space.

Quantum algorithms. By a polynomial-time *quantum algorithm* (or QPT algorithm) we mean a polynomial-time uniform family of quantum circuits given by $\mathcal{C} = \bigcup_{n\in\mathbb{N}} C_n$, where each circuit $C \in \mathcal{C}$ is described by a sequence of unitary gates and measurements. Similarly, we also define (classical) probabilistic polynomial-time (PPT) algorithms. A quantum algorithm may, in general, receive (mixed) quantum states as inputs and produce (mixed) quantum states as outputs. Occasionally, we restrict QPT algorithms implicitly. For example, if we write $\Pr[\mathcal{A}(1^\lambda) = 1]$ for a QPT algorithm $\mathcal{A}$, it is implicit that $\mathcal{A}$ is a QPT algorithm that outputs a single classical bit.

We extend the notion of QPT algorithms to CPTP channels via the following definition.

Definition 1 (Efficient CPTP maps)**.** *A family of* CPTP *maps* $\{\Phi_\lambda : L(\mathcal{H}_{A_\lambda}) \to L(\mathcal{H}_{B_\lambda})\}_{\lambda\in\mathbb{N}}$ *is called efficient, if there exists a polynomial-time uniformly generated family of circuits* $\{C_\lambda\}_{\lambda\in\mathbb{N}}$ *acting on the Hilbert space* $\mathcal{H}_{A_\lambda} \otimes \mathcal{H}_{B_\lambda} \otimes \mathcal{H}_{C_\lambda}$ *such that, for all* $\lambda \in \mathbb{N}$ *and for all* $\varrho \in \mathcal{H}_{A_\lambda}$*,*

$$\Phi_\lambda(\varrho_\lambda) = \mathrm{Tr}_{A_\lambda C_\lambda}[C_\lambda(\varrho_\lambda \otimes |0\rangle\langle 0|_{B_\lambda C_\lambda})].$$

Definition 2 (Indistinguishability of ensembles of random variables)**.** *Let* $\lambda \in N$ *be a parameter. We say that two ensembles of random variables* $X = \{X_\lambda\}$ *and* $Y = \{Y_\lambda\}$ *are computationally indistinguishable, denoted by* $X \approx_c Y$*, if for all* QPT *distinguishers* $\mathcal{D}$ *which output a single bit, it holds that*

$$\left|\Pr[\mathcal{D}(1^\lambda, X_\lambda) = 1] - \Pr[\mathcal{D}(1^\lambda, Y_\lambda) = 1]\right| \leq \mathrm{negl}(\lambda)\,.$$

Definition 3 (Indistinguishability of ensembles of quantum states, [Wat06])**.** *Let* $p : \mathbb{N} \to \mathbb{N}$ *be a polynomially bounded function, and let* ϱ_λ *and* σ_λ *be* $p(\lambda)$*-qubit quantum states. We say that* $\{\varrho_\lambda\}_{\lambda\in\mathbb{N}}$ *and* $\{\sigma_\lambda\}_{\lambda\in\mathbb{N}}$ *are quantum computationally indistinguishable ensembles of quantum states, denoted by* $\varrho_\lambda \approx_c \sigma_\lambda$*, if, for any* QPT *distinguisher* $\mathcal{D}$ *with single-bit output, any polynomially bounded* $q : \mathbb{N} \to \mathbb{N}$*, any family of* $q(\lambda)$*-qubit auxiliary states* $\{\nu_\lambda\}_{\lambda\in\mathbb{N}}$*, and every* $\lambda \in \mathbb{N}$*,*

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

Lemma 1 ("Almost As Good As New" Lemma, [Aar16])**.** *Let* $\varrho \in \mathcal{D}(\mathcal{H})$ *be a density matrix over a Hilbert space* $\mathcal{H}$*. Let* U *be an arbitrary unitary and let* $(\mathbf{\Pi}_0, \mathbf{\Pi}_1 = \mathbb{1} - \mathbf{\Pi}_0)$ *be projectors acting on* $\mathcal{H} \otimes \mathcal{H}_{\mathsf{aux}}$*. We interpret* $(U, \mathbf{\Pi}_0, \mathbf{\Pi}_1)$ *as a measurement performed by appending an ancillary system in the state* $|0\rangle\langle 0|_{\mathsf{aux}}$*, applying the unitary* U *and subsequently performing the two-outcome measurement* $\{\mathbf{\Pi}_0, \mathbf{\Pi}_1\}$ *on the larger*

system. Suppose that the outcome corresponding to $\mathbf{\Pi}_0$ *occurs with probability* $1-\varepsilon$*, for some* $\varepsilon \in [0,1]$*. In other words, it holds that* $\mathrm{Tr}[\mathbf{\Pi}_0(U\varrho \otimes |0\rangle\langle 0|_{\mathsf{aux}} U^\dagger)] = 1-\varepsilon$*. Then,*

$$\|\widetilde{\varrho} - \varrho\|_{\mathrm{tr}} \leq \sqrt{\varepsilon},$$

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

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

We also use the following lemma on the closeness to ideal states:

Lemma 2 ([Unr13], Lemma 10)**.** *Let* $\mathbf{\Pi}$ *be an arbitrary projector and let* $|\psi\rangle$ *be a normalized pure state such that* $\|\mathbf{\Pi}|\psi\rangle\|^2 = 1-\varepsilon$*, for some* $\varepsilon \geq 0$*. Then, there exists a (pure) ideal state,*

$$|\bar{\psi}\rangle = \frac{\mathbf{\Pi}|\psi\rangle}{\|\mathbf{\Pi}|\psi\rangle\|},$$

with the property that

$$\||\psi\rangle\langle\psi| - |\bar{\psi}\rangle\langle\bar{\psi}|\|_{\mathrm{tr}} \leq \sqrt{\varepsilon} \quad \text{and} \quad |\bar{\psi}\rangle \in \mathrm{im}(\mathbf{\Pi}).$$

In other words, the state $|\bar{\psi}\rangle$ *is within trace distance* $\varepsilon > 0$ *of the state* $|\psi\rangle$ *and lies in the image of* $\mathbf{\Pi}$*.*

We also use the following elementary lemma.

Lemma 3 ([CMP20], Lemma 23)**.** *Let* $\varrho, \sigma \in \mathcal{D}(\mathcal{H})$ *be two states with the property that* $\|\varrho - \sigma\|_{\mathrm{tr}} \leq \varepsilon$*, for some* $\varepsilon \geq 0$*. Let* $\mathbf{\Pi}$ *be an arbitrary matrix acting on* $\mathcal{H}$ *such that* $0 \leq \mathbf{\Pi} \leq \mathbb{1}$*. Then,*

$$|\mathrm{Tr}[\mathbf{\Pi}\varrho] - \mathrm{Tr}[\mathbf{\Pi}\sigma]| \leq \varepsilon.$$

2.2 Classical and quantum entropies

Classical entropies. Let X be a random variable with an arbitrary distribution $P_{\mathcal{X}}$ over an alphabet $\mathcal{X}$. The *min-entropy* of X, denoted by $H_{\min}(X)$, is defined by the following quantity

$$H_{\min}(X) = -\log\left(\max_{x\in\mathcal{X}} \Pr_{X\sim P_{\mathcal{X}}}[X = x]\right).$$

The *conditional min-entropy* of X conditioned on a correlated random variable Y is defined by

$$H_{\min}(X|Y) = -\log\left(\mathop{\mathbb{E}}_{y\leftarrow Y}\left[\max_{x\in\mathcal{X}} \Pr_{X\sim P_{\mathcal{X}}}[X = x|Y = y]\right]\right).$$

Lemma 4 (Leftover Hash Lemma, [HILL88])**.** *Let* $n, m \in \mathbb{N}$ *and* $q \geq 2$ *a prime. Let* P *be a distribution over* $\mathbb{Z}_q^m$ *and suppose that* $H_{\min}(X) \geq n\log q + 2\log(1/\varepsilon) + O(1)$ *for* $\varepsilon > 0$*, where* X *denotes a random variable with distribution* P*. Then, the following two distributions are within total variance distance* ε*:*

$$(\mathbf{A}, \mathbf{A}\cdot\mathbf{x} \pmod q) \quad \approx_\varepsilon \quad (\mathbf{A}, \mathbf{u}) : \qquad \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m},\ \mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^n.$$

Quantum entropies.

Definition 4 (Quantum min-entropy)**.** *Let A and B be two quantum systems and let $\varrho_{AB} \in \mathcal{S}_{\leq}(\mathcal{H}_{AB})$ be any bipartite state. The min-entropy of A conditioned on B of the state ϱ_{AB} is defined as*

$$H_{\min}(A \mid B)_{\varrho} = \max_{\sigma \in S_{\leq}(\mathcal{H}_B)} \sup \left\{ \lambda \in \mathbb{R} \, : \, \varrho_{AB} \leq 2^{-\lambda} \mathbb{1}_A \otimes \sigma_B \right\}.$$

Definition 5 (Smooth quantum min-entropy)**.** *Let A and B be quantum systems and let $\varrho_{AB} \in \mathcal{S}_{\leq}(\mathcal{H}_{AB})$. Let $\varepsilon \geq 0$. We define the ε-smooth quantum min-entropy of A conditioned on B of ϱ_{AB} as*

$$H^{\varepsilon}_{\min}(A \mid B)_{\varrho} \; = \sup_{\substack{\tilde{\varrho}_{AB} \\ P(\tilde{\varrho}_{AB}, \varrho_{AB}) \leq \varepsilon}} H_{\min}(A \mid B)_{\tilde{\varrho}}.$$

The conditional min-entropy of a CQ state ϱ_{XB} captures the difficulty of guessing the content of a classical register X given quantum side information B. This motivates the following definition.

Definition 6 (Guessing probability)**.** *Let $\varrho_{XB} \in \mathcal{D}(\mathcal{H}_X \otimes \mathcal{H}_B)$ be a CQ state, where X is a classical register over an alphabet $\mathcal{X}$ and B is a quantum system. Then, the guessing probability of X given B is defined as*

$$p_{\text{guess}}(X|B)_{\varrho} = \sup_{M^x_B} \sum_{x \in \mathcal{X}} \Pr[X = x]_{\varrho} \cdot \operatorname{Tr}\left[M^x_B \varrho_B M^{x\dagger}_B \right].$$

The following operational meaning of min-entropy is due to Koenig, Renner and Schaffner [KRS09].

Theorem 1 ([KRS09], Theorem 1)**.** *Let $\varrho_{XB} \in \mathcal{D}(\mathcal{H}_X \otimes \mathcal{H}_B)$ be a CQ state, where X is a classical register over an alphabet $\mathcal{X}$ and B is a quantum system. Then, it holds that*

$$H_{\min}(X \mid B)_{\varrho} = -\log \left(p_{\text{guess}}(X|B)_{\varrho} \right).$$

2.3 Fourier analysis

Let $q \geq 2$ be an integer modulus and let $m \in \mathbb{N}$. The *q-ary (discrete) Fourier transform* takes as input a function $f : \mathbb{Z}^m \to \mathbb{C}$ and produces a function $\hat{f} : \mathbb{Z}^m_q \to \mathbb{C}$ (the Fourier transform of f) defined by

$$\hat{f}(\mathbf{y}) = \sum_{\mathbf{x} \in \mathbb{Z}^m} f(\mathbf{x}) \cdot e^{\frac{2\pi i}{q} \langle \mathbf{y}, \mathbf{x} \rangle}.$$

For brevity, we oftentimes write $\omega_q = e^{\frac{2\pi i}{q}} \in \mathbb{C}$ to denote the primitive q-th root of unity. The m-qudit *q-ary quantum Fourier transform* over the ring $\mathbb{Z}^m_q$ is defined by the operation,

$$\mathsf{FT}_q : \quad |\mathbf{x}\rangle \quad \mapsto \quad \sqrt{q^{-m}} \sum_{\mathbf{y} \in \mathbb{Z}^n_q} e^{\frac{2\pi i}{q} \langle \mathbf{y}, \mathbf{x} \rangle} |\mathbf{y}\rangle, \qquad \forall \mathbf{x} \in \mathbb{Z}^m_q.$$

It is well known that the q-ary quantum Fourier transform can be efficiently performed on a quantum computer for any modulus $q \geq 2$ [HH00]. Note the quantum Fourier transform of a normalized quantum state

$$|\Psi\rangle = \sum_{\mathbf{x} \in \mathbb{Z}^m} f(\mathbf{x}) |\mathbf{x}\rangle \qquad \text{with} \qquad \sum_{\mathbf{x} \in \mathbb{Z}^m} |f(\mathbf{x})|^2 = 1,$$

for a function $f : \mathbb{Z}^m \to \mathbb{C}$, results in the state (the Fourier transform of $|\Psi\rangle$) given by

$$\begin{aligned}\mathsf{FT}_q |\Psi\rangle &= \sqrt{q^{-m}} \sum_{\mathbf{y}\in\mathbb{Z}_q^n} \Big(\sum_{\mathbf{x}\in\mathbb{Z}^m} f(\mathbf{x}) \cdot e^{\frac{2\pi i}{q}\langle \mathbf{y},\mathbf{x}\rangle} \Big) |\mathbf{y}\rangle \\ &= \sqrt{q^{-m}} \sum_{\mathbf{y}\in\mathbb{Z}_q^n} \hat{f}(\mathbf{y}) |\mathbf{y}\rangle .\end{aligned}$$

Notice that the Fourier transform of $|\Psi\rangle$ is *unitary* if $\mathrm{supp}(f) \subseteq \mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$. We frequently make use of the following standard identity for Fourier characters.

Lemma 5 (Orthogonality of Fourier characters)**.** *Let $q \geq 2$ be any integer modulus and let $\omega_q = e^{\frac{2\pi i}{q}} \in \mathbb{C}$ denote the primitive q-th root of unity. Then, for arbitrary $x, y \in \mathbb{Z}_q$:*

$$\sum_{v\in\mathbb{Z}_q} \omega_q^{v\cdot x} \omega_q^{-v\cdot y} = q\ \delta_{x,y}.$$

2.4 Generalized Pauli operators

Definition 7 (Generalized Pauli operators)**.** *Let $q \geq 2$ be an integer modulus and $\omega_q = e^{2\pi i/q}$ be the primitive q-th root of unity. The generalized q-ary Pauli operators $\{\mathbf{X}_q^b\}_{b\in\mathbb{Z}_q}$ and $\{\mathbf{Z}_q^b\}_{b\in\mathbb{Z}_q}$ are given by*

$$\begin{aligned}\mathbf{X}_q^b &= \sum_{a\in\mathbb{Z}_q} |a + b \pmod q\rangle \langle a| , \quad \text{and} \\ \mathbf{Z}_q^b &= \sum_{a\in\mathbb{Z}_q} \omega_q^{a\cdot b} |a\rangle \langle a| .\end{aligned}$$

For $\mathbf{b} = (b_1, \ldots, b_m) \in \mathbb{Z}_q^m$, we use the notation $\mathbf{X}_q^{\mathbf{b}} = \mathbf{X}_q^{b_1} \otimes \cdots \otimes \mathbf{X}_q^{b_m}$ and $\mathbf{Z}_q^{\mathbf{b}} = \mathbf{Z}_q^{b_1} \otimes \cdots \otimes \mathbf{Z}_q^{b_m}$.

Lemma 6. *Let $q \geq 2$ be an integer modulus. Then, for all $b \in \mathbb{Z}_q$, it holds that*

$$\begin{aligned}\mathbf{Z}_q^b &= \mathsf{FT}_q\, \mathbf{X}_q^b\, \mathsf{FT}_q^\dagger \\ \mathbf{X}_q^b &= \mathsf{FT}_q^\dagger\, \mathbf{Z}_q^b\, \mathsf{FT}_q.\end{aligned}$$

Proof. It suffices to show the first identity only as the second identity follows by conjugation with FT_q. Using the orthogonality of Fourier characters over $\mathbb{Z}_q$ (Lemma 5), we find that

$$\begin{aligned}\mathbf{Z}_q^b &= \sum_{x\in\mathbb{Z}_q} \omega_q^{x\cdot b} |x\rangle\langle x| \\ &= \sum_{x,y'\in\mathbb{Z}_q} \omega_q^{x\cdot b} \left(\frac{1}{q} \sum_{a\in\mathbb{Z}_q} \omega_q^{x\cdot a} \omega_q^{-a\cdot y'} \right) |x\rangle\langle y'| \\ &= \frac{1}{q} \sum_{x,y\in\mathbb{Z}_q} \sum_{x',y'\in\mathbb{Z}_q} \sum_{a\in\mathbb{Z}_q} \omega_q^{x\cdot y} \omega_q^{-x'\cdot y'} \langle y| a + b \pmod q\rangle \cdot \langle a|x'\rangle\ |x\rangle\langle y'| \\ &= \frac{1}{q} \left(\sum_{x,y\in\mathbb{Z}_q} \omega_q^{x\cdot y} |x\rangle\langle y| \right) \sum_{a\in\mathbb{Z}_q} |a + b \pmod q\rangle \langle a| \left(\sum_{x',y'\in\mathbb{Z}_q} \omega_q^{-x'\cdot y'} |\mathbf{x}'\rangle\langle \mathbf{y}'| \right) \\ &= \mathsf{FT}_q\, \mathbf{X}_q^b\, \mathsf{FT}_q^\dagger.\end{aligned}$$

□

Definition 8 (Pauli-**Z** dephasing channel)**.** *Let $q \geq 2$ be an integer modulus and let $m \in \mathbb{N}$. Let $\mathbf{p}$ be a probability distribution over $\mathbb{Z}_q^m$. Then, the Pauli-**Z** dephasing channel with respect to $\mathbf{p}$ is defined as*

$$\mathcal{Z}_{\mathbf{p}}(\varrho) = \sum_{\mathbf{z}\in\mathbb{Z}_q^m} p_{\mathbf{z}}\, \mathbf{Z}_q^{\mathbf{z}} \varrho \mathbf{Z}_q^{-\mathbf{z}}, \qquad \forall \varrho \in L((\mathbb{C}^q)^{\otimes m}).$$

*We use $\mathcal{Z}$ to denote the uniform Pauli-**Z** channel for which $\mathbf{p}$ is the uniform distribution over $\mathbb{Z}_q^m$.*

The following lemma shows that the uniform Pauli-**Z** channel on input ϱ returns a diagonal state which consists of diagonal elements of ϱ encoded in the standard basis.

Lemma 7. *Let $q \geq 2$ be a modulus and $m \in \mathbb{N}$. Then, the uniform Pauli-**Z** dephasing channel satsifies,*

$$\mathcal{Z}(\varrho) = q^{-m} \sum_{\mathbf{z}\in\mathbb{Z}_q^m} \mathbf{Z}_q^{\mathbf{z}} \varrho \mathbf{Z}_q^{-\mathbf{z}} = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \mathrm{Tr}[|\mathbf{x}\rangle\langle\mathbf{x}|\, \varrho]\ |\mathbf{x}\rangle\langle\mathbf{x}|\,, \qquad \forall \varrho \in L((\mathbb{C}^q)^{\otimes m}).$$

Proof. Suppose that the state ϱ has the following form in the standard basis,

$$\varrho = \sum_{\mathbf{x},\mathbf{y}\in\mathbb{Z}_q^m} \alpha_{\mathbf{x},\mathbf{y}}\, |\mathbf{x}\rangle\langle\mathbf{y}| \ \in L((\mathbb{C}^q)^{\otimes m}).$$

Using the orthogonality of Fourier characters over $\mathbb{Z}_q$ (Lemma 5), we obtain

$$\begin{aligned}
\mathcal{Z}(\varrho) &= q^{-m} \sum_{\mathbf{z}\in\mathbb{Z}_q^m} \mathbf{Z}_q^{\mathbf{z}} \varrho \mathbf{Z}_q^{-\mathbf{z}} \\
&= q^{-m} \sum_{\mathbf{z}\in\mathbb{Z}_q^m} \sum_{\mathbf{x},\mathbf{y}\in\mathbb{Z}_q^m} \alpha_{\mathbf{x},\mathbf{y}}\, \mathbf{Z}_q^{\mathbf{z}} |\mathbf{x}\rangle\langle\mathbf{y}| \mathbf{Z}_q^{-\mathbf{z}} \\
&= \sum_{\mathbf{x},\mathbf{y}\in\mathbb{Z}_q^m} \alpha_{\mathbf{x},\mathbf{y}} \left(q^{-m} \sum_{\mathbf{z}\in\mathbb{Z}_q^m} \omega_q^{\langle \mathbf{x},\mathbf{z}\rangle} \omega_q^{-\langle \mathbf{y},\mathbf{z}\rangle} \right) |\mathbf{x}\rangle\langle\mathbf{y}| \\
&= \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \alpha_{\mathbf{x},\mathbf{x}}\, |\mathbf{x}\rangle\langle\mathbf{x}| \\
&= \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \mathrm{Tr}[|\mathbf{x}\rangle\langle\mathbf{x}|\, \varrho]\, |\mathbf{x}\rangle\langle\mathbf{x}|\,.
\end{aligned}$$

□

2.5 Lattices and the Gaussian mass

A *lattice* $\Lambda \subset \mathbb{R}^m$ is a discrete subgroup of $\mathbb{R}^m$. To avoid handling matters of precision, we will only consider integer lattices $\Lambda \subseteq \mathbb{Z}^m$ throughout this work. The *dual* of a lattice $\Lambda \subset \mathbb{R}^m$, denoted by Λ^*, is the lattice of all vectors $y \in \mathbb{R}^m$ that satisfy $\langle \mathbf{y}, \mathbf{x}\rangle \in \mathbb{Z}$, for all vectors $\mathbf{x} \in \Lambda$. In other words, we define

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

Given a lattice $\Lambda \subset \mathbb{R}^m$ and a vector $\mathbf{t} \in \mathbb{R}^m$, we define the coset with respect to $\mathbf{t}$ as the lattice shift $\Lambda - \mathbf{t} = \{\mathbf{x} \in \mathbb{R}^m : \mathbf{x} + \mathbf{t} \in \Lambda\}$. Note that many different shifts $\mathbf{t}$ can define the same coset.

The *Gaussian measure* ϱ_σ with parameter $\sigma > 0$ is defined as the function

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

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

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

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

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

We make use of the following tail bound for the Gaussian mass of a lattice [Ban93, Lemma 1.5 (ii)].

Lemma 8. *For any m-dimensional lattice Λ and shift $\mathbf{t} \in \Lambda$ and for all $r > 0$, $c \geq (2\pi)^{-\frac{1}{2}}$ it holds that*

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

where $B^m(\mathbf{0}, s) = \{\mathbf{x} \in \mathbb{R}^m : \|\mathbf{x}\|_2 \leq s\}$ denotes the m-dimensional ball of radius $s > 0$.

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

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

Note that $\Lambda_q^\perp(\mathbf{A})$ contains $q\mathbb{Z}^m$; in particular, it contains the identity $\mathbf{0} \in \mathbb{Z}^m$. For any *syndrome* $\mathbf{y} \in \mathbb{Z}_q^n$ in the column span of $\mathbf{A}$, we also consider the lattice coset $\Lambda_q^\mathbf{y}(\mathbf{A})$ given by

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

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

The second lattice is the lattice generated by $\mathbf{A}^T$ and is defined by

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

The q-ary lattices $\Lambda_q(\mathbf{A})$ and $\Lambda_q^\perp(\mathbf{A})$ are dual to each other (up to scaling). Specifically, we have

$$q \cdot \Lambda_q^\perp(\mathbf{A})^* = \Lambda_q(\mathbf{A}) \quad \text{and} \quad q \cdot \Lambda_q(\mathbf{A})^* = \Lambda_q^\perp(\mathbf{A}).$$

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

Lemma 9 ([GPV07], Lemma 5.1)**.** *Let $n \in \mathbb{N}$ and let $q \geq 2$ be a prime modulus with $m \geq 2n \log q$. Then, for all but a q^{-n} fraction of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, the subset-sums of the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$. In other words, a uniformly random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ is full-rank with overwhelming probability.*

Lemma 10 ([GPV07], Corollary 5.4)**.** *Let $n \in \mathbb{N}$ and $q \geq 2$ be a prime with $m \geq 2n \log q$. Then, for all but a $2q^{-n}$ fraction of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\sigma = \omega(\sqrt{\log m})$, the distribution of the syndrome $\mathbf{A} \cdot \mathbf{e} = \mathbf{u} \pmod q$ is within negligible total variation distance of the uniform distribution over $\mathbb{Z}_q^n$, where $\mathbf{e} \sim D_{\mathbb{Z}^m,\sigma}$.*

The following lemma is a consequence of [MR04, Lemma 4.4] and [GPV07, Lemma 5.3].

Lemma 11. *Let $n \in \mathbb{N}$ and let $q \geq 2$ be a prime modulus with $m \geq 2n \log q$. Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$. Then, for any $\sigma = \omega(\sqrt{\log m})$ and for any syndrome $\mathbf{y} \in \mathbb{Z}_q^n$:*

$$\Pr_{\mathbf{x} \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}} \left[\|\mathbf{x}\| \geq \sqrt{m}\sigma \right] \leq \mathrm{negl}(n).$$

Definition 9 (Periodic Gaussian)**.** *Let $m \in \mathbb{N}$, let $q \geq 2$ be a modulus and let $\sigma > 0$. The q-periodic Gaussian $\varrho_{\sigma,q}$ function is the periodic continuation of the Gaussian measure ϱ_σ, where*

$$\varrho_{\sigma,q}(\mathbf{x}) = \varrho_\sigma(\mathbf{x} + q\mathbb{Z}^m), \quad \forall \mathbf{x} \in \mathbb{R}^m.$$

For any function $f : \mathbb{Z}^m \to \mathbb{C}$ and lattice $\Lambda \subseteq \mathbb{Z}^m$, the well-known *Poisson summation formula* states that $f(\Lambda) = \det(\Lambda^*)\hat{f}(\Lambda^*)$. We use the following Gaussian variant of the formula [Bra18, Corollary 2.14].

Lemma 12 (Poisson summation formula)**.** *Let $q \geq 2$ be a prime modulus and let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be any matrix whose columns generate $\mathbb{Z}_q^n$. Let $\mathbf{v}, \mathbf{w} \in \mathbb{Z}_q^m$ and $\sigma > 0$ be arbitrary. Then, it holds that*

$$\sum_{\mathbf{x} \in \Lambda_q^{\mathbf{v}}(\mathbf{A})} \varrho_\sigma(\mathbf{x}) \cdot e^{-\frac{2\pi i}{q}\langle \mathbf{w}, \mathbf{x} \rangle} = \frac{\sigma^m}{q^n} \cdot \sum_{\mathbf{y} \in \mathbb{Z}_q^n} \varrho_{q/\sigma,q}(\mathbf{w} + \mathbf{y}\mathbf{A}) \cdot e^{\frac{2\pi i}{q}\langle \mathbf{y}, \mathbf{v} \rangle}.$$

For $\mathbf{x} \in \mathbb{Z}^m$, let $[\mathbf{x}]_q$ denote the unique representative $\bar{\mathbf{x}} \in \mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$ such that $\mathbf{x} \equiv \bar{\mathbf{x}} \pmod{q}$. The following lemma due to Brakerski [Bra18] says that, whenever σ is much smaller than the modulus q, the periodic Gaussian $\varrho_{\sigma,q}$ is close to the non-periodic (but truncated) Gaussian.

Lemma 13 ([Bra18], Lemma 2.6)**.** *Let $q \geq 2$, $\mathbf{x} \in \mathbb{Z}^m$ such that $\|[\mathbf{x}]_q\| < q/4$ and $\sigma > 0$. Then,*

$$1 \leq \frac{\varrho_{\sigma,q}(\mathbf{x})}{\varrho_\sigma([\mathbf{x}]_q)} \leq 1 + 2^{-(\frac{1}{2}(q/\sigma)^2 - m)}.$$

A simple consequence of the tail bound in Lemma 8 is that the discrete Gaussian $D_{\mathbb{Z}^m,\sigma}$ distribution is essentially only supported on the finite set $\{\mathbf{x} \in \mathbb{Z}^m : \|\mathbf{x}\|_\infty \leq \sigma\sqrt{m}\}$, which suggests the use of *truncation*. Given a modulus $q \geq 2$ and $\sigma > 0$, we define the *truncated* discrete Gaussian distribution $D_{\mathbb{Z}_q^m,\sigma}$ over the finite set $\mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$ with support $\{\mathbf{x} \in \mathbb{Z}_q^m : \|\mathbf{x}\|_\infty \leq \sigma\sqrt{m}\}$ as the density

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

We define the analogous *periodic* discrete Gaussian distribution $D_{\mathbb{Z}_q^m,\sigma,q}$ as

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

Lemma 14. *Let $m \in \mathbb{N}$, $q \geq 2$ a modulus and let $\sigma \in (0, q/\sqrt{8m})$. Consider the quantum states,*

$$|\psi\rangle = \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m,\sigma}(\mathbf{x})}\,|\mathbf{x}\rangle \qquad \text{and} \qquad |\phi\rangle = \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m,\sigma,q}(\mathbf{x})}\,|\mathbf{x}\rangle\,.$$

Then, it holds that

$$\| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_{\mathrm{tr}} \leq \sqrt{1 - \left(1 + 2^{-(\frac{1}{2}(q/\sigma)^2 - m)}\right)^{-1}}.$$

Proof. We first bound the Hellinger distance,

$$H^2(D_{\mathbb{Z}_q^m,\sigma}, D_{\mathbb{Z}_q^m,\sigma,q}) = 1 - \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m,\sigma}(\mathbf{x}) \cdot D_{\mathbb{Z}_q^m,\sigma,q}(\mathbf{x})}. \tag{5}$$

To this end, we define two normalization factors

$$Z_\sigma = \sum_{\mathbf{y}\in\mathbb{Z}_q^m, \|\mathbf{y}\|_\infty \le \sqrt{m}\sigma} \varrho_\sigma(\mathbf{y}) \qquad \text{and} \qquad Z_{\sigma,q} = \sum_{\mathbf{y}\in\mathbb{Z}_q^m, \|\mathbf{y}\|_\infty \le \sqrt{m}\sigma} \varrho_{\sigma,q}(\mathbf{y}). \tag{6}$$

From Lemma 13, it follows for any $\mathbf{x} \in \mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$ with $\|\mathbf{x}\| < q/4$ that

$$\varrho^2_{\sigma,q}(\mathbf{x}) \cdot \left(1 + 2^{-(\frac{1}{2}(q/\sigma)^2 - m)}\right)^{-1} \le \varrho_\sigma(\mathbf{x}) \cdot \varrho_{\sigma,q}(\mathbf{x}). \tag{7}$$

Recall also that the truncated discrete Gaussian is supported on the finite set

$$\mathrm{supp}(D_{\mathbb{Z}_q^m,\sigma}) = \{\mathbf{x} \in \mathbb{Z}_q^m : \|\mathbf{x}\|_\infty \le \sqrt{m}\sigma\}.$$

Plugging in Eq. (7), we can bound the Hellinger distance as follows:

$$\begin{aligned}
H^2(D_{\mathbb{Z}_q^m,\sigma}, D_{\mathbb{Z}_q^m,\sigma,q}) &= 1 - \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m,\sigma}(\mathbf{x}) \cdot D_{\mathbb{Z}_q^m,\sigma,q}(\mathbf{x})} \\
&= 1 - \sqrt{Z_\sigma^{-1} \cdot Z_{\sigma,q}^{-1}} \sum_{\mathbf{x}\in\mathbb{Z}_q^m, \|\mathbf{x}\|_\infty \le \sqrt{m}\sigma} \sqrt{\varrho_\sigma(\mathbf{x}) \cdot \varrho_{\sigma,q}(\mathbf{x})} \\
&\le 1 - \sqrt{\frac{Z_\sigma^{-1} \cdot Z_{\sigma,q}^{-1}}{1 + 2^{-(\frac{1}{2}(q/\sigma)^2 - m)}}} \sum_{\mathbf{x}\in\mathbb{Z}_q^m, \|\mathbf{x}\|_\infty \le \sqrt{m}\sigma} \varrho_{\sigma,q}(\mathbf{x}) \\
&\le 1 - \left(1 + 2^{-(\frac{1}{2}(q/\sigma)^2 - m)}\right)^{-1/2}.
\end{aligned}$$

Therefore, it holds that

$$\begin{aligned}
\| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_{\mathrm{tr}} &\le \sqrt{1 - (1 - H^2(D_{\mathbb{Z}_q^m,\sigma}, D_{\mathbb{Z}_q^m,\sigma,q}))^2} \\
&\le \sqrt{1 - \left(1 + 2^{-(\frac{1}{2}(q/\sigma)^2 - m)}\right)^{-1}}.
\end{aligned}$$

□

The following result allows us to bound the total variation distance between a truncated discrete Gaussian $D_{\mathbb{Z}_q^m,\sigma}$ and its perturbation by a fixed vector $\mathbf{e}_0 \in \mathbb{Z}^m$.

Lemma 15 ([BCM+21], Lemma 2.4)**.** *Let $q \ge 2$ be a modulus, $m \in \mathbb{N}$ and $\sigma > 0$. Then, for any $\mathbf{e}_0 \in \mathbb{Z}^m$,*

$$\|D_{\mathbb{Z}_q^m,\sigma} - (D_{\mathbb{Z}_q^m,\sigma} + \mathbf{e}_0)\|_{\mathsf{TV}} \le 2 \cdot \left(1 - e^{\frac{-2\pi\sqrt{m}\|\mathbf{e}_0\|}{\sigma}}\right).$$

2.6 Cryptography

In this section, we review several definitions in cryptography.

Public-key encryption.

Definition 10 (Public-key encryption)**.** *A public-key encryption* (PKE) *scheme* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ *with plaintext space* $\mathcal{M}$ *is a triple of* PPT *algorithms consisting of a key generation algorithm* KeyGen*, an encryption algorithm* Enc*, and a decryption algorithm* Dec*.*

$\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: *takes as input the parameter* 1^λ *and outputs a public key* pk *and secret key* sk*.*

$\mathsf{Enc}(\mathsf{pk}, m) \to \mathsf{CT}$: *takes as input the public key* pk *and a plaintext* $m \in \mathcal{M}$*, and outputs a ciphertext* CT*.*

$\mathsf{Dec}(\mathsf{sk}, \mathsf{CT}) \to m'$ **or** $\perp$: *takes as input the secret key* sk *and ciphertext* CT*, and outputs* $m' \in \mathcal{M}$ *or* $\perp$*.*

Definition 11 (Correctness of PKE)**.** *For any* $\lambda \in \mathbb{N}$*, and for any* $m \in \mathcal{M}$*:*

$$\Pr\left[\mathsf{Dec}(\mathsf{sk}, \mathsf{CT}) \neq m \;\middle|\; \begin{smallmatrix}(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{pk}, m)\end{smallmatrix}\right] \leq \mathrm{negl}(\lambda).$$

Definition 12 (IND-CPA security)**.** *Let* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ *be a* PKE *scheme and let* $\mathcal{A}$ *be a* QPT *adversary. We define the security experiment* $\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(b)$ *between* $\mathcal{A}$ *and a challenger as follows:*

1. *The challenger generates a pair* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$*, and sends* pk *to* $\mathcal{A}$*.*
2. $\mathcal{A}$ *sends a plaintext pair* $(m_0, m_1) \in \mathcal{M} \times \mathcal{M}$ *to the challenger.*
3. *The challenger computes* $\mathsf{CT}_b \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b)$*, and sends* CT_b *to* $\mathcal{A}$*.*
4. $\mathcal{A}$ *outputs a bit* $b' \in \{0,1\}$*, which is also the output of the experiment.*

We say that the scheme Σ *is* IND-CPA*-secure if, for any* QPT *adversary* $\mathcal{A}$*, it holds that*

$$\mathsf{Adv}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

2.7 The Short Integer Solution problem

The (inhomogenous) SIS problem was introduced by Ajtaj [Ajt96] in his seminal work on average-case lattice problems. The problem is defined as follows.

Definition 13 (Inhomogenous SIS problem, [Ajt96])**.** *Let* $n, m \in \mathbb{N}$ *be integers, let* $q \geq 2$ *be a modulus and let* $\beta > 0$ *be a parameter. The Inhomogenous Short Integer Solution problem* (ISIS) *problem is to find a short solution* $\mathbf{x} \in \mathbb{Z}^m$ *with* $\|\mathbf{x}\|_2 \leq \beta$ *such that* $\mathbf{A} \cdot \mathbf{x} = \pmod q$ *given as input a tuple* $(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \mathbf{y} \xleftarrow{\$} \mathbb{Z}_q^n)$*. The Short Integer Solution* (SIS) *problem is a homogenous variant of* ISIS *with input* $(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \mathbf{0} \in \mathbb{Z}_q^n)$*.*

Micciancio and Regev [MR07] showed that the average-case SIS problem is as hard as approximating worst-case lattice problems to within small factors. Gentry, Peikert and Vaikuntanathan [GPV07] subsequently gave an improved reduction showing that, for any $m = \mathrm{poly}(n)$, $\beta = \mathrm{poly}(n)$, and prime $q \geq \beta \cdot \omega(\sqrt{n \log q})$, the average-case problems $\mathsf{SIS}_{n,m,q,\beta}$ and $\mathsf{ISIS}_{n,m,q,\beta}$ are as hard as approximating the shortest independent vector problem (SIVP) problem in the worst case to within a factor $\gamma = \beta \cdot \tilde{O}(\sqrt{n})$.

2.8 The Learning with Errors problem

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

Definition 14 ("Search" LWE, [Reg05])**.** *Let $n, m \in \mathbb{N}$ be integers, let $q \geq 2$ be a modulus and let $\alpha \in (0,1)$ be a parameter. The Learning with Errors (LWE) problem is to find a secret vector $\mathbf{s}$ given as input a sample $(\mathbf{A}, \mathbf{sA} + \mathbf{e} \pmod q)$ from the distribution $\mathsf{LWE}^m_{n,q,\alpha q}$, where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ are uniformly random, and where $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ is sampled from the discrete Gaussian distribution.*

Definition 15 ("Decisional" LWE, [Reg05])**.** *Let $n, m \in \mathbb{N}$ be integers, let $q \geq 2$ be a modulus and let $\alpha \in (0,1)$ be a parameter. The "decision" Learning with Errors (DLWE) problem is to distinguish between*

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

where $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ is uniformly random and where $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ is a discrete Gaussian noise vector.

As shown in [Reg05], the $\mathsf{LWE}^m_{n,q,\alpha q}$ problem with parameter $\alpha q \geq 2\sqrt{n}$ is at least as hard as approximating the shortest independent vector problem (SIVP) to within a factor of $\gamma = \widetilde{O}(n/\alpha)$ in worst case lattices of dimension n. In this work we assume the subexponential hardness of $\mathsf{LWE}^m_{n,q,\alpha q}$ which relies on the worst case hardness of approximating short vector problems in lattices to within a subexponential factor.

3 Primal and Dual Gaussian States

Our Dual-Regev-type encryption schemes with certified deletion in Section 7 and Section 9 rely on two types of Gaussian superpositions, which we call *primal* and *dual* Gaussian states. The former (i.e., primal) state corresponds to a quantum superposition of LWE samples with respect to a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, and (up to a phase) can be thought of as a superposition of Gaussian balls around random lattice vectors in $\Lambda_q(\mathbf{A})$. The latter (i.e., dual) state corresponds to a Gaussian superposition over a particular coset,

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

of the q-ary lattice $\Lambda_q^{\perp}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\cdot\mathbf{x} = \mathbf{0} \pmod q\}$ defined in Section 2.5.

Our terminology regarding which state is primal and which state is dual is completely arbitrary. In fact, the q-ary lattices $\Lambda_q(\mathbf{A})$ and $\Lambda_q^{\perp}(\mathbf{A})$ are both dual to each other (up to scaling), and satisfy

$$q\cdot\Lambda_q^{\perp}(\mathbf{A})^* = \Lambda_q(\mathbf{A}) \quad \text{and} \quad q\cdot\Lambda_q(\mathbf{A})^* = \Lambda_q^{\perp}(\mathbf{A}).$$

We choose to refer to the quantum superposition of LWE samples as the *primal* Gaussian state because it corresponds directly to the ciphertexts of our encryption scheme, whereas the *dual* Fourier mode is only used in order to prove deletion. Gaussian superpositions first appeared in Regev's quantum reduction from worst-case lattice problems to LWE, and have also been used by Stehlé et al. [SSTX09] who gave a quantum reduction from SIS to LWE. Roberts [Rob19] and Kitagawa et al. [KNY21] used similar (dual) Gaussian states to construct quantum money and secure software leasing schemes. Various other forms of superpositions of LWE samples have been considered by Grilo, Kerenidis and Zijlstra [GKZ19] in the context of quantum learning theory and by Alagic, Jeffery, Ozols and Poremba [AJOP20], as well as by Chen, Liu and Zhandry [CLZ21], in the context of quantum cryptanalysis of LWE-based cryptosystems.

We define primal and dual Gaussian states as follows.

Definition 16 (Gaussian states)**.** *Let $m \in \mathbb{N}$, $q \geq 2$ be an integer modulus and $\sigma > 0$. Then,*

- *(primal Gaussian state:) for all $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and $\mathbf{y} \in \mathbb{Z}_q^m$, we let*

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

- *(dual Gaussian state:) for all $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and $\mathbf{y} \in \mathbb{Z}_q^m$, we let*

$$|\hat{\psi}_{\mathbf{A},\mathbf{y}}\rangle \;=\; \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x})\, |\mathbf{x}\rangle\,.$$

For simplicity, we oftentimes drop the subscript on $\mathbf{A}$ and write $|\psi_{\mathbf{y}}\rangle$ and $|\hat{\psi}_{\mathbf{y}}\rangle$, respectively.

3.1 Transference lemma

The following lemma states that, up to negligible trace distance, the primal and dual Gaussian states in Definition 16 are related via the q-ary quantum Fourier transform.

Lemma 16 (Transference lemma)**.** *Let $m \in \mathbb{N}$, $q \geq 2$ be a prime modulus and let $\sigma \in (\sqrt{8m}, q/\sqrt{8m})$. Let $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$ and let $\mathbf{y} \in \mathbb{Z}_q^n$ be arbitrary. Then, up to negligible trace distance, the primal and dual Gaussian states are related via the quantum Fourier transform:*

$$\begin{aligned}
\mathsf{FT}_q\,|\psi_{\mathbf{y}}\rangle \quad &\approx_\varepsilon \quad |\hat{\psi}_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x})\, |\mathbf{x}\rangle\,; \\
\mathsf{FT}_q^\dagger\,|\hat{\psi}_{\mathbf{y}}\rangle \quad &\approx_\varepsilon \quad |\psi_{\mathbf{y}}\rangle = \sum_{\mathbf{s}\in\mathbb{Z}_q^n} \sum_{\mathbf{e}\in\mathbb{Z}_q^m} \varrho_{q/\sigma}(\mathbf{e})\, \omega_q^{-\langle \mathbf{s},\mathbf{y}\rangle}\, |\mathbf{s}\mathbf{A}+\mathbf{e} \pmod q\rangle\,,
\end{aligned}$$

where $\varepsilon : \mathbb{N} \to \mathbb{R}^+$ is a negligible function in the parameter $m \in \mathbb{N}$.

Proof. Let $\mathbf{y} \in \mathbb{Z}_q^n$ be an arbitrary vector and recall that the dual Gaussian coset $|\hat{\psi}_{\mathbf{y}}\rangle$ is given by

$$|\hat{\psi}_{\mathbf{y}}\rangle \;=\; \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x})\, |\mathbf{x}\rangle\,. \tag{8}$$

We denote by $\Lambda_q^{\mathbf{y}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\mathbf{x} = \mathbf{y} \pmod q\}$ be the associated coset of the lattice $\Lambda_q^\perp(\mathbf{A})$. Consider now the Gaussian superposition over the entire lattice coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ formally defined by

$$|\hat{\phi}_{\mathbf{y}}\rangle \;=\; \sum_{\mathbf{x}\in\Lambda_q^{\mathbf{y}}(\mathbf{A})} \varrho_\sigma(\mathbf{x})\, |\mathbf{x}\rangle\,. \tag{9}$$

Since $\sigma < q/\sqrt{8m}$, it follows from the tail bound in Lemma 11 that the state in (8) is within negligible trace distance of the state in Eq. (9). Applying the (inverse) quantum Fourier transform, we get

$$|\phi_{\mathbf{y}}\rangle \stackrel{\text{def}}{=} \mathsf{FT}_q^\dagger\,|\hat{\phi}_{\mathbf{y}}\rangle = \sum_{\mathbf{z}\in\mathbb{Z}_q^m} \Big(\sum_{\mathbf{x}\in\Lambda_q^{\mathbf{y}}(\mathbf{A})} \varrho_\sigma(\mathbf{x}) \cdot \omega_q^{-\langle \mathbf{x},\mathbf{z}\rangle} \Big)\, |\mathbf{z}\rangle\,. \tag{10}$$

From the Poisson summation formula (Lemma 12) and a subsequent change of variables, it follows that

$$\begin{aligned}|\phi_{\mathbf{y}}\rangle &= \sum_{\mathbf{z}\in\mathbb{Z}_q^m}\Big(\sum_{\mathbf{s}\in\mathbb{Z}_q^n}\varrho_{q/\sigma,q}(\mathbf{z}+\mathbf{sA})\cdot\omega_q^{\langle\mathbf{s},\mathbf{y}\rangle}\Big)|\mathbf{z}\rangle \\ &= \sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m}\varrho_{q/\sigma,q}(\mathbf{e})\cdot\omega_q^{-\langle\mathbf{s},\mathbf{y}\rangle}\,|\mathbf{sA}+\mathbf{e}\ (\mathrm{mod}\ q)\rangle\,. \end{aligned} \tag{11}$$

Because $\sigma > \sqrt{8m}$ it follows from Lemma 14 that there exists

$$\kappa(m) = \sqrt{1-(1+2^{-3m})^{-1}} \geq 0$$

such that

$$|\phi_{\mathbf{y}}\rangle \approx_{\kappa} \sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m}\varrho_{q/\sigma}(\mathbf{e})\cdot\omega_q^{-\langle\mathbf{s},\mathbf{y}\rangle}\,|\mathbf{sA}+\mathbf{e}\ (\mathrm{mod}\ q)\rangle\,. \tag{12}$$

Putting everything together, it follows from the triangle inequality that

$$\mathsf{FT}_q^{\dagger}\,|\hat{\psi}_{\mathbf{y}}\rangle \quad \approx_{\varepsilon} \quad |\psi_{\mathbf{y}}\rangle = \sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m}\varrho_{q/\sigma}(\mathbf{e})\,\omega_q^{-\langle\mathbf{s},\mathbf{y}\rangle}\,|\mathbf{sA}+\mathbf{e}\rangle\,,$$

where $\varepsilon(m) = \mathrm{negl}(m) + \kappa(m)$. Using that $\sqrt{1-1/(1+x)} \leq \sqrt{x}$ for all $x > 0$, we have

$$\begin{aligned}\varepsilon(m) &= \mathrm{negl}(m) + \sqrt{1-(1+2^{-3m})^{-1}} \\ &\leq \mathrm{negl}(m) + 2^{-\frac{3m}{2}}.\end{aligned}$$

Thus, we have that $\varepsilon(m) \leq \mathrm{negl}(m)$. This proves the claim. □

Corollary 1. *Let $m \in \mathbb{N}$, $q \geq 2$ be a prime and $\sigma \in (\sqrt{8m}, q/\sqrt{8m})$. Let $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$ and let $\mathbf{y} \in \mathbb{Z}_q^n$ be arbitrary. Then, there exists a negligible function $\varepsilon(m)$ such that*

$$\mathsf{FT}_q\mathbf{X}_q^{\mathbf{v}}\,|\psi_{\mathbf{y}}\rangle \quad \approx_{\varepsilon} \quad \mathbf{Z}_q^{\mathbf{v}}\,|\hat{\psi}_{\mathbf{y}}\rangle\,, \qquad \forall \mathbf{v} \in \mathbb{Z}_q^m.$$

Proof. From Lemma 6 it follows that $\mathsf{FT}_q\mathbf{X}_q^{\mathbf{v}} = \mathbf{Z}_q^{\mathbf{v}}\mathsf{FT}_q$, for all $\mathbf{v} \in \mathbb{Z}_q^m$. Moreover, Lemma 16 implies that $\mathsf{FT}_q\,|\psi_{\mathbf{y}}\rangle$ is within negligible trace distance of $|\hat{\psi}_{\mathbf{y}}\rangle$. This proves the claim. □

3.2 Efficient state preparation

In this section, we give two algorithms that prepare the *primal* and *dual* Gaussian states from Definition 16. We remark that Gaussian superpositions over $\mathbb{Z}_q^m$ with parameter $\sigma = \Omega(\sqrt{m})$ can be efficiently implemented using standard quantum state preparation techniques, for example using *rejection sampling* and the *Grover-Rudolph algorithm*. We refer to [GR02, Reg05, Bra18, BCM+21]) for a reference.

Our first algorithm (see Algorithm 1 in Figure 2) prepares the dual Gaussian state from Definition 16 with respect to an input matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and parameter $\sigma = \Omega(\sqrt{m})$, and is defined as follows.

Our second algorithm (see Algorithm 2 in Figure 3) prepares the primal Gaussian state with respect to an input matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and parameter $\sigma = O(\frac{q}{\sqrt{m}})$. Here, in order for Lemma 16 to apply, it is crucial that the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$. Fortunately, it follows from Lemma 9 that a uniformly random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ satisfies this property with overwhelming probability.

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

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

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

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

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

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

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

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

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

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

Figure 2: Quantum algorithm which takes as input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and a width parameter $\sigma = \Omega(\sqrt{m})$, and outputs the dual Gaussian state in Definition 16.

3.3 Invariance under Pauli-Z dephasing

In this section, we prove a surprising property about the dual Gaussian state from Definition 16. We prove Theorem 2, which says that the Pauli-$\mathbf{Z}$ dephasing channel with respect to the LWE distribution leaves the dual Gaussian state approximately invariant.

Theorem 2. *Let $n, m \in \mathbb{N}$ be integers and let $q \geq 2$ be a prime modulus, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\sigma \in (\sqrt{8m}, q/\sqrt{8m})$ be a function of λ. Let $\mathbf{y} \in \mathbb{Z}_q^n$ be any vector and $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ be any matrix whose columns generate $\mathbb{Z}_q^n$, and let $|\hat{\psi}_{\mathbf{y}}\rangle$ be the dual Gaussian state,*

$$|\hat{\psi}_{\mathbf{y}}\rangle \;=\; \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle .$$

Let $\mathcal{Z}_{\mathsf{LWE}^m_{n,q,\alpha q}}$ be the Pauli-$\mathbf{Z}$ dephasing channel with respect to the $\mathsf{LWE}^m_{n,q,\alpha q}$ distribution for $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and a noise ratio $\alpha \in (0,1)$ with relative noise magnitude $\frac{q/\sigma}{\alpha q} = \lambda^{\omega(1)}$, i.e.

$$\mathcal{Z}_{\mathsf{LWE}^m_{n,q,\alpha q}}(\varrho) = \sum_{\mathbf{s}_0\in\mathbb{Z}_q^n} \sum_{\mathbf{e}_0\in\mathbb{Z}_q^m} q^{-n} D_{\mathbb{Z}_q^m,\alpha q}(\mathbf{e}_0)\, \mathbf{Z}_q^{\mathbf{s}_0\cdot\mathbf{A}+\mathbf{e}_0}\, \varrho\, \mathbf{Z}_q^{-(\mathbf{s}_0\cdot\mathbf{A}+\mathbf{e}_0)}, \qquad \forall \varrho \in L((\mathbb{C}^q)^{\otimes m}).$$

Algorithm 2: $\mathsf{GenPrimal}(\mathbf{A}, \sigma)$

Input: Matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ whose columns generate $\mathbb{Z}_q^n$, and a parameter $\sigma = O(\frac{q}{\sqrt{m}})$.

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

1 Run $\mathsf{GenDual}(\mathbf{A}, \sigma)$, resulting in the state

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

2 Apply the quantum Fourier transform FT_q to system X.

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

Figure 3: Quantum algorithm which takes as input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and a parameter $\sigma = O(\frac{q}{\sqrt{m}})$, and outputs the primal Gaussian state in Definition 16.

Then, there exists a negligible function $\varepsilon(\lambda)$ such that

$$\mathcal{Z}_{\mathsf{LWE}^m_{n,q,\alpha q}}(|\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|) \;\approx_\varepsilon\; |\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|\,.$$

In other words, the Pauli-$\mathbf{Z}$ dephasing channel with respect to the LWE *distribution leaves the dual Gaussian state approximately invariant.*

Proof. Let $\mathbf{y} \in \mathbb{Z}_q^n$ be an arbitrary vector and recall that the dual Gaussian state $|\hat{\psi}_{\mathbf{y}}\rangle$ is given by

$$|\hat{\psi}_{\mathbf{y}}\rangle \;=\; \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m\\ \mathbf{A}\mathbf{x}=\mathbf{y}\ (\mathrm{mod}\ q)}} \varrho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle\,. \tag{13}$$

Consider a sample $\mathbf{b} = \mathbf{s}_0\cdot\mathbf{A} + \mathbf{e}_0 \ (\mathrm{mod}\ q)) \sim \mathsf{LWE}^m_{n,q,\alpha q}$ with $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ and $\mathbf{e}_0 \sim D_{\mathbb{Z}_q^m,\alpha q}$. Because $\sigma \in (\sqrt{8m}, q/\sqrt{8m})$ and $\frac{q/\sigma}{\alpha q} = \lambda^{\omega(1)}$, there exist negligible $\eta(\lambda)$ and $\kappa(\lambda)$ such that

$$\begin{aligned}
\mathbf{Z}_q^{s_0\cdot\mathbf{A}+\mathbf{e}_0}\,|\hat{\psi}_{\mathbf{y}}\rangle \;&=\; \mathsf{FT}_q\,\mathbf{X}_q^{s_0\cdot\mathbf{A}+\mathbf{e}_0}\,\mathsf{FT}_q^\dagger\,|\hat{\psi}_{\mathbf{y}}\rangle && \text{(Lemma 6)}\\
&\approx_\eta\; \mathsf{FT}_q\,\mathbf{X}_q^{s_0\cdot\mathbf{A}+\mathbf{e}_0}\,|\psi_{\mathbf{y}}\rangle && \text{(Lemma 16)}\\
&\approx_\kappa\; \omega_q^{\langle\mathbf{s}_0,\mathbf{y}\rangle}\mathsf{FT}_q\,|\psi_{\mathbf{y}}\rangle && \text{(Lemma 15)}\\
&\approx_\eta\; \omega_q^{\langle\mathbf{s}_0,\mathbf{y}\rangle}\,|\hat{\psi}_{\mathbf{y}}\rangle\,. && \text{(Lemma 16)}
\end{aligned}$$

Here, $|\psi_{\mathbf{y}}\rangle$ is the primal Gaussian state given by

$$|\psi_{\mathbf{y}}\rangle = \sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m} \varrho_{q/\sigma}(\mathbf{e})\,\omega_q^{-\langle\mathbf{s},\mathbf{y}\rangle}\,|\mathbf{s}\mathbf{A}+\mathbf{e}\ (\mathrm{mod}\ q)\rangle\,.$$

In other words, $|\hat{\psi}_{\mathbf{y}}\rangle$ in Eq. (13) is an approximate eigenvector of the generalized Pauli operator $\mathbf{Z}_q^{s_0\cdot\mathbf{A}+\mathbf{e}_0}$ with respect to the same matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$. Note that we can simply discard $\omega_q^{\langle\mathbf{s}_0,\mathbf{y}\rangle} \in \mathbb{C}$ because it serves

as a global phase. Hence, there exists a negligible function $\varepsilon(\lambda)$ such that

$$\begin{aligned}\mathcal{Z}_{\mathsf{LWE}^m_{n,q,\alpha q}}(|\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|) &= \sum_{\mathbf{s}_0\in\mathbb{Z}_q^n}\sum_{\mathbf{e}_0\in\mathbb{Z}_q^m} q^{-n} D_{\mathbb{Z}_q^m,\alpha q}(\mathbf{e}_0)\, \mathbf{Z}_q^{s_0\cdot\mathbf{A}+\mathbf{e}_0}\, |\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|\, \mathbf{Z}_q^{-(s_0\cdot\mathbf{A}+\mathbf{e}_0)} \\ &\approx_\varepsilon \left(\sum_{\mathbf{s}_0\in\mathbb{Z}_q^n} q^{-n}\right)\cdot\left(\sum_{\mathbf{e}_0\in\mathbb{Z}_q^m} D_{\mathbb{Z}_q^m,\alpha q}(\mathbf{e}_0)\right)\, |\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}| \\ &= |\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|\,.\end{aligned}$$

□

4 Uncertainty Relation for Fourier Basis Projections

In this section, we prove an entropic uncertainty relation with respect to so-called Fourier basis projections. Informally, we say that a projector $\widehat{\mathbf{\Pi}}$ is a *Fourier basis projection*, if $\widehat{\mathbf{\Pi}}$ corresponds to a projector (onto a subset of $\mathbb{Z}_q^m$) which is conjugated by the q-ary Fourier transform FT_q. Notice that the deletion procedures of our encryption schemes with certified deletion in Section 7 and Section 9 require a Fourier basis projection onto a small set of solutions to the (inhomogenous) short integer solution (ISIS) problem. Another example can be found in the work of Aaronson and Christiano [AC12] who used Hadamard basis projections (a special case of the q-ary Fourier transform) onto small hidden subspaces to verify quantum money states.

Our uncertainty relation captures the following intuitive property: any system which passes a Fourier basis projection onto a small subset of $\mathbb{Z}_q^m$ (say, with high probability) must necessarily be *unentangled* with any auxiliary system. We formalize this statement using the (smooth) quantum min-entropy (Definition 5).

4.1 Fourier basis projections

Definition 17 (Fourier basis projection)**.** *Let $m\in\mathbb{N}$ and let $q\geq 2$ be an integer modulus. Let $\mathcal{S}\subseteq\mathbb{Z}_q^m$ be an arbitrary set and let $\mathbf{\Pi}_{\mathcal{S}}$ be the associated projector onto $\mathcal{S}$, where*

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

Then, we define the associated Fourier basis projection onto $\mathcal{S}$ as the projector

$$\widehat{\mathbf{\Pi}}_{\mathcal{S}} = \mathsf{FT}_q^\dagger \mathbf{\Pi}_{\mathcal{S}} \mathsf{FT}_q.$$

4.2 Uncertainty relation

In this section, our main result is the following.

Theorem 3 (Uncertainty relation for Fourier basis projections)**.** *Let $m\in\mathbb{N}$, $q\geq 2$ be a modulus, $\{|\psi^{\mathbf{x}}\rangle\}_{\mathbf{x}\in\mathbb{Z}_q^m}$ be any family of normalized auxiliary states, and let $|\psi\rangle_{AB}$ be any state of the form*

$$|\psi\rangle_{AB} = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \alpha_{\mathbf{x}}\, |\mathbf{x}\rangle_A \otimes |\psi^{\mathbf{x}}\rangle_B \qquad s.t. \qquad \sum_{\mathbf{x}\in\mathbb{Z}_q^m} |\alpha_{\mathbf{x}}|^2 = 1.$$

Let $\mathcal{S}\subseteq\mathbb{Z}_q^m$ be an arbitrary set and define the following projectors onto system A,

$$\mathbf{\Pi}_{\mathcal{S}} = \sum_{\mathbf{x}\in\mathcal{S}} |\mathbf{x}\rangle\langle\mathbf{x}| \qquad \textit{and} \qquad \widehat{\mathbf{\Pi}}_{\mathcal{S}} = \mathsf{FT}_q^\dagger \mathbf{\Pi}_{\mathcal{S}} \mathsf{FT}_q.$$

Suppose that $\|(\widehat{\mathbf{\Pi}}_{\mathcal{S}} \otimes \mathbb{1}_B)\,|\psi\rangle_{AB}\,\|^2 = 1-\varepsilon$, *for some* $\varepsilon \geq 0$. *Then, it holds that*

$$H_{\min}^{\sqrt{\varepsilon}}(X|B)_{\varrho} \;\geq\; m \cdot \log q - 2 \cdot \log |\mathcal{S}|.$$

Here, ϱ_{XB} *results from a computational basis measurement of system* A *of the state* $|\psi\rangle\langle\psi|_{AB}$, *i.e.*

$$\varrho_{XB} = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} |\mathbf{x}\rangle\langle\mathbf{x}|_X \otimes \mathrm{tr}_A\left[(|\mathbf{x}\rangle\langle\mathbf{x}|_A \otimes \mathbb{1}_B)|\psi\rangle\langle\psi|_{AB}\right].$$

Proof. Suppose that $|\psi\rangle_{AB}$ satisfies $\|(\widehat{\mathbf{\Pi}}_{\mathcal{S}} \otimes \mathbb{1}_B)\,|\psi\rangle_{AB}\,\|^2 = 1-\varepsilon$, for some $\varepsilon \geq 0$. From Lemma 2, it follows that there exists an ideal pure state,

$$|\bar{\psi}\rangle_{AB} = \frac{(\widehat{\mathbf{\Pi}}_{\mathcal{S}} \otimes \mathbb{1}_B)\,|\psi\rangle_{AB}}{\|(\widehat{\mathbf{\Pi}}_{\mathcal{S}} \otimes \mathbb{1}_B)\,|\psi\rangle_{AB}\,\|} = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \bar{\alpha}_{\mathbf{x}}\,|\mathbf{x}\rangle_A \otimes |\psi^{\mathbf{x}}\rangle_B \qquad \text{s.t.} \qquad \sum_{\mathbf{x}\in\mathbb{Z}_q^m} |\bar{\alpha}_{\mathbf{x}}|^2 = 1,$$

with the property that

$$\||\psi\rangle\langle\psi| - |\bar{\psi}\rangle\langle\bar{\psi}|\|_{\mathrm{tr}} \leq \sqrt{\varepsilon} \qquad \text{and} \qquad |\bar{\psi}\rangle \in \mathrm{im}(\widehat{\mathbf{\Pi}}_{\mathcal{S}} \otimes \mathbb{1}_B).$$

Because $|\bar{\psi}\rangle_{AB}$ lies in the image of the projector $\widehat{\mathbf{\Pi}}_{\mathcal{S}} \otimes \mathbb{1}_B$, we have

$$|\bar{\psi}\rangle_{AB} = (\widehat{\mathbf{\Pi}}_{\mathcal{S}} \otimes \mathbb{1}_B)\,|\bar{\psi}\rangle_{AB} = \sum_{\mathbf{x}'\in\mathbb{Z}_q^m} \sum_{\mathbf{s}\in\mathcal{S}} \bar{\alpha}_{\mathbf{x}'} \cdot \omega_q^{\langle\mathbf{x},\mathbf{s}\rangle}\omega_q^{-\langle\mathbf{x}',\mathbf{s}\rangle}\,|\mathbf{x}\rangle_A \otimes |\psi^{\mathbf{x}'}\rangle_B\,.$$

Let us now analyze the ideal state $\bar{\varrho}_{XB}$ which results from a computational basis measurement of system A of the state $|\bar{\psi}\rangle\langle\bar{\psi}|_{AB}$. In other words, we consider the CQ state given by

$$\bar{\varrho}_{XB} = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} |\mathbf{x}\rangle\langle\mathbf{x}|_X \otimes \mathrm{tr}_A\left[(|\mathbf{x}\rangle\langle\mathbf{x}|_A \otimes \mathbb{1}_B)|\bar{\psi}\rangle\langle\bar{\psi}|_{AB}\right].$$

By the definition of the guessing probability in Definition 6, we have

$$\begin{aligned}
p_{\mathrm{guess}}(X|B)_{\bar{\varrho}} &= \sup_{M_B^x} \sum_{x\in\mathbb{Z}_q^m} \left\|(|\mathbf{x}\rangle\langle\mathbf{x}|_A \otimes \boldsymbol{M}_B^{\mathbf{x}})\,|\bar{\psi}\rangle_{AB}\right\|^2 \\
&= \sup_{M_B^x} \sum_{x\in\mathbb{Z}_q^m} q^{-2m} \left\| \sum_{\mathbf{x}'\in\mathbb{Z}_q^m} \sum_{\mathbf{s}\in\mathcal{S}} \bar{\alpha}_{\mathbf{x}} \cdot \omega_q^{\langle\mathbf{x},\mathbf{s}\rangle}\omega_q^{-\langle\mathbf{x}',\mathbf{s}\rangle}\,|\mathbf{x}\rangle_A \otimes \boldsymbol{M}_B^{\mathbf{x}}\,|\psi^{\mathbf{x}'}\rangle_B \right\|^2 \\
&= \sup_{M_B^x} \sum_{x\in\mathbb{Z}_q^m} q^{-2m} \left\| \sum_{\mathbf{x}'\in\mathbb{Z}_q^m} \bar{\alpha}_{\mathbf{x}'} \cdot \Big(\sum_{\mathbf{s}\in\mathcal{S}} \omega_q^{\langle\mathbf{x},\mathbf{s}\rangle}\omega_q^{-\langle\mathbf{x}',\mathbf{s}\rangle} \Big)\,|\mathbf{x}\rangle_A \otimes \boldsymbol{M}_B^{\mathbf{x}}\,|\psi^{\mathbf{x}'}\rangle_B \right\|^2 .
\end{aligned}$$

Using the Cauchy-Schwarz-inequality, we find that for any $\mathbf{x} \in \mathbb{Z}_q^m$:

$$\begin{aligned}
&\left\| \sum_{\mathbf{x}' \in \mathbb{Z}_q^m} \bar{\alpha}_{\mathbf{x}'} \cdot \Big(\sum_{\mathbf{s} \in \mathcal{S}} \omega_q^{\langle \mathbf{x},\mathbf{s} \rangle} \omega_q^{-\langle \mathbf{x}',\mathbf{s} \rangle} \Big) \, |\mathbf{x}\rangle_A \otimes \boldsymbol{M}_B^{\mathbf{x}} \, |\psi^{\mathbf{x}'}\rangle_B \right\| \\
&\leq \sqrt{\sum_{\mathbf{x}' \in \mathbb{Z}_q^m} \left| \bar{\alpha}_{\mathbf{x}'} \cdot \Big(\sum_{\mathbf{s} \in \mathcal{S}} \omega_q^{\langle \mathbf{x},\mathbf{s} \rangle} \omega_q^{-\langle \mathbf{x}',\mathbf{s} \rangle} \Big) \right|^2} \cdot \sqrt{\sum_{\mathbf{x}' \in \mathbb{Z}_q^m} \left\| \, |\mathbf{x}\rangle_A \otimes \boldsymbol{M}_B^{\mathbf{x}} \, |\psi^{\mathbf{x}'}\rangle_B \right\|^2} \\
&\leq \sqrt{|\mathcal{S}|^2 \sum_{\mathbf{x}' \in \mathbb{Z}_q^m} \left| \bar{\alpha}_{\mathbf{x}'} \right|^2} \cdot \sqrt{\sum_{\mathbf{x}' \in \mathbb{Z}_q^m} \left\| \, |\mathbf{x}\rangle_A \otimes \boldsymbol{M}_B^{\mathbf{x}} \, |\psi^{\mathbf{x}'}\rangle_B \right\|^2} \\
&= |\mathcal{S}| \cdot \sqrt{\sum_{\mathbf{x}' \in \mathbb{Z}_q^m} \left\| \boldsymbol{M}_B^{\mathbf{x}} \, |\psi^{\mathbf{x}'}\rangle_B \right\|^2}. \qquad (14)
\end{aligned}$$

Using the inequality in (14), we can now bound the guessing probability as follows:

$$\begin{aligned}
p_{\text{guess}}(X|B)_{\bar{\varrho}} \;&\leq\; \frac{|\mathcal{S}|^2}{q^{2m}} \cdot \sup_{\boldsymbol{M}_B^{\mathbf{x}}} \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \sum_{\mathbf{x}' \in \mathbb{Z}_q^m} \left\| \boldsymbol{M}_B^{\mathbf{x}} \, |\psi^{\mathbf{x}'}\rangle_B \right\|^2 \\
&= \frac{|\mathcal{S}|^2}{q^{2m}} \cdot \sum_{\mathbf{x}' \in \mathbb{Z}_q^m} \sup_{\boldsymbol{M}_B^{\mathbf{x}}} \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \left\| \boldsymbol{M}_B^{\mathbf{x}} \, |\psi^{\mathbf{x}'}\rangle_B \right\|^2 \\
&= \frac{|\mathcal{S}|^2}{q^m}. \qquad \text{(since } \sum_{\mathbf{x}} \boldsymbol{M}_B^{\mathbf{x}} = \mathbb{1}\text{)}
\end{aligned}$$

Because the *purified distance* is bounded above by the *trace distance*, it follows that

$$P(\varrho_{XB}, \bar{\varrho}_{XB}) \;\leq\; \|\varrho_{XB} - \bar{\varrho}_{XB}\|_{\text{tr}} \;\leq\; \||\psi\rangle\langle\psi| - |\bar{\psi}\rangle\langle\bar{\psi}|\|_{\text{tr}} \;\leq\; \sqrt{\varepsilon}.$$

Therefore, by the definition of (smooth) min-entropy (see Definition 5), we have

$$H_{\min}(X \,|\, B)_{\bar{\varrho}} \;\leq\; \sup_{\substack{\sigma_{XB} \\ P(\sigma_{XB}, \varrho_{XB}) \leq \sqrt{\varepsilon}}} H_{\min}(X \,|\, B)_{\sigma} \;=\; H_{\min}^{\sqrt{\varepsilon}}(X \,|\, B)_{\varrho}. \qquad (15)$$

Putting everything together, it follows from (15) and Theorem 1 that

$$\begin{aligned}
H_{\min}^{\sqrt{\varepsilon}}(X \,|\, B)_{\varrho} \;&\geq\; H_{\min}(X \,|\, B)_{\bar{\varrho}} \\
&= -\log \left(p_{\text{guess}}(X|B)_{\bar{\varrho}} \right) \\
&\geq\; m \cdot \log q - 2 \cdot \log |\mathcal{S}|.
\end{aligned}$$

This proves the claim. □

5 Gaussian-Collapsing Hash Functions

Unruh [Unr15] introduced the notion of collapsing hash functions in his seminal work on computationally binding quantum commitments. This property is captured by the following definition.

Definition 18 (Collapsing hash function, [Unr15])**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. A hash function family $\mathcal{H} = \{H_\lambda\}_{\lambda\in\mathbb{N}}$ is called collapsing if, for every* QPT *adversary $\mathcal{A}$,*

$$|\Pr[\mathsf{CollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{CollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

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

1. *The challenger samples a random hash function $h \xleftarrow{\$} H_\lambda$, and sends a description of h to $\mathcal{A}$.*
2. *$\mathcal{A}$ responds with a (classical) string $y \in \{0,1\}^{n(\lambda)}$ and an $m(\lambda)$-qubit quantum state in system X.*
3. *The challenger coherently computes h (into an auxiliary system Y) given the state in system X, and then performs a two-outcome measurement on Y indicating whether the output of h equals y. If h does not equal y the challenger aborts and outputs $\perp$.*
4. *If $b = 0$, the challenger does nothing. Else, if $b = 1$, the challenge measures the $m(\lambda)$-qubit system X in the computational basis. Finally, the challenger returns the state in system X to $\mathcal{A}$.*
5. *$\mathcal{A}$ returns a bit b', which we define as the output of the experiment.*

Motivated by the properties of the dual Gaussian state from Definition 16, we consider a special class of hash functions which are *collapsing* with respect to Gaussian superpositions. Informally, we say that a hash function h is *Gaussian-collapsing* if it is computationally difficult to distinguish between a Gaussian superposition of pre-images and a single (measured) Gaussian pre-image (of h). We formalize this below.

Definition 19 (Gaussian-collapsing hash function)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter, $m(\lambda), n(\lambda) \in \mathbb{N}$ and let $q(\lambda) \geq 2$ be a modulus. Let $\sigma > 0$. A hash function family $\mathcal{H} = \{H_\lambda\}_{\lambda\in\mathbb{N}}$ with domain $\mathcal{X} = \mathbb{Z}_q^m$ and range $\mathcal{Y} = \mathbb{Z}_q^n$ is called σ-Gaussian-collapsing if, for every* QPT *adversary $\mathcal{A}$,*

$$|\Pr[\mathsf{GaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{GaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

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

1. *The challenger samples a random hash function $h \xleftarrow{\$} H_\lambda$ and prepares the quantum state*

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

2. *The challenger measures system Y in the computational basis, resulting in the state*

$$|\hat{\psi}_{\mathbf{y}}\rangle_{XY} = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m: \\ h(\mathbf{x})=\mathbf{y}}} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle_X \otimes |\mathbf{y}\rangle_Y .$$

3. *If $b = 0$, the challenger does nothing. Else, if $b = 1$, the challenger measures system X of the quantum state $|\hat{\psi}_{\mathbf{y}}\rangle$ in the computational basis. Finally, the challenger sends the outcome state in systems X to $\mathcal{A}$, together with the string $\mathbf{y} \in \mathbb{Z}_q^n$ and a classical description of the hash function h.*
4. *$\mathcal{A}$ returns a bit b', which we define as the output of the experiment.*

The following follows immediately from the definition of Gaussian-collapsing hash functions, and the fact that the dual Gaussian state can be efficiently prepared using Algorithm 1.

Claim 1. *Let $\mathcal{H} = \{H_\lambda\}_{\lambda\in\mathbb{N}}$ be a hash function family with domain $\mathcal{X} = \mathbb{Z}_q^m$ and range $\mathcal{Y} = \mathbb{Z}_q^n$, where $m(\lambda), n(\lambda) \in \mathbb{N}$. If $\mathcal{H}$ is collapsing, then $\mathcal{H}$ is also σ-Gaussian-collapsing, for any $\sigma = \Omega(\sqrt{m})$.*

5.1 Ajtaj's hash function

Liu and Zhandry [LZ19] implicitly showed that the *Ajtaj* hash function $h_{\mathbf{A}}(\mathbf{x}) = \mathbf{A}\mathbf{x} \pmod q$ is collapsing – and thus *Gaussian-collapsing* – via the notion of *lossy functions* and by assuming the superpolynomial hardness of (decisional) LWE. In this section, we give a simple and direct proof that the Ajtaj hash function is Gaussian-collapsing assuming (decisional) LWE, which might be of independent interest.

Theorem 4. *Let $n, m \in \mathbb{N}$ be integers and let $q \geq 2$ be a prime modulus, each parameterized by $\lambda \in \mathbb{N}$. Let $\sigma \in (\sqrt{8m}, q/\sqrt{8m})$ be a function of λ. Then, the Ajtaj hash function family $\mathcal{H} = \{H_\lambda\}_{\lambda\in\mathbb{N}}$ with*

$$H_\lambda = \left\{ h_{\mathbf{A}} : \mathbb{Z}_q^m \to \mathbb{Z}_q^n \ \textit{s.t.}\ h_{\mathbf{A}}(\mathbf{x}) = \mathbf{A}\cdot\mathbf{x} \pmod q;\ \mathbf{A} \in \mathbb{Z}_q^{n\times m} \right\}$$

is σ-Gaussian-collapsing assuming the quantum hardness of the decisional $\mathsf{LWE}^m_{n,q,\alpha q}$ problem, for any parameter $\alpha \in (0,1)$ with relative noise magnitude $\frac{q/\sigma}{\alpha q} = \lambda^{\omega(1)}$.

Proof. Let $\mathcal{A}$ denote the QPT adversary in the experiment $\mathsf{GaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(b)$ with $b \in \{0,1\}$. To prove the claim, we give a reduction from the decisional $\mathsf{LWE}^m_{n,q,\alpha q}$ assumption. We are given as input a sample $(\mathbf{A},\mathbf{b})$ with $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$, where $\mathbf{b} = \mathbf{s}_0\cdot\mathbf{A} + \mathbf{e}_0 \pmod q)$ is either a sample from the LWE distribution with $\mathbf{s}_0 \xleftarrow{\$} \mathbb{Z}_q^n$ and $\mathbf{e}_0 \sim D_{\mathbb{Z}^m,\alpha q}$, or where $\mathbf{b}$ is a uniformly random string $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$.

Consider the distinguisher $\mathcal{D}$ that acts as follows on input 1^λ and $(\mathbf{A},\mathbf{b})$:

1. $\mathcal{D}$ prepares a bipartite quantum state on systems X and Y with
$$|\hat{\psi}\rangle_{XY} = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \varrho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle_X \otimes |\mathbf{A}\cdot\mathbf{x} \pmod q\rangle_Y .$$
2. $\mathcal{D}$ measures system Y in the computational basis, resulting in the state
$$|\hat{\psi}_{\mathbf{y}}\rangle_{XY} = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{x}=\mathbf{y}}} \varrho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle_X \otimes |\mathbf{y}\rangle_Y .$$
3. $\mathcal{D}$ applies the generalized Pauli-$\mathbf{Z}$ operator $\mathbf{Z}_q^{\mathbf{b}}$ on system X, resulting in the state
$$(\mathbf{Z}_q^{\mathbf{b}}\otimes\mathbb{1}_Y)\,|\hat{\psi}_{\mathbf{y}}\rangle_{XY} = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{x}=\mathbf{y}}} \varrho_\sigma(\mathbf{x})\left(\mathbf{Z}_q^{\mathbf{b}}\,|\mathbf{x}\rangle_X\right)\otimes|\mathbf{y}\rangle_Y .$$
4. $\mathcal{D}$ runs the adversary $\mathcal{A}$ on input system X and classical descriptions of $\mathbf{A}\in\mathbb{Z}_q^{n\times m}$ and $\mathbf{y}\in\mathbb{Z}_q^n$.
5. $\mathcal{D}$ outputs whatever bit $b'\in\{0,1\}$ the adversary $\mathcal{A}$ outputs.

Suppose that, for every $\lambda\in\mathbb{N}$, there exists a polynomial $p(\lambda)$ such that

$$|\Pr[\mathsf{GaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(0)=1] - \Pr[\mathsf{GaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(1)=1]| \geq \frac{1}{p(\lambda)}.$$

We now show that this implies that $\mathcal{D}$ succeeds at the decisional $\mathsf{LWE}^m_{n,q,\alpha q}$ experiment with advantage at least $1/p(\lambda) - \mathrm{negl}(\lambda)$. We distinguish between the following two cases.

If $(\mathbf{A}, \mathbf{b})$ is a sample from the LWE distribution with $\mathbf{b} = \mathbf{s}_0 \cdot \mathbf{A} + \mathbf{e}_0 \pmod q$), then the adversary $\mathcal{A}$ receives as input the following quantum state in system X:

$$\mathcal{Z}_{\mathsf{LWE}^m_{n,q,\alpha q}}(|\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|_X) = \sum_{\mathbf{s}_0\in\mathbb{Z}_q^n}\sum_{\mathbf{e}_0\in\mathbb{Z}^m} q^{-n} D_{\mathbb{Z}^m,\alpha q}(\mathbf{e}_0)\ \mathbf{Z}_q^{s_0\cdot\mathbf{A}+\mathbf{e}_0}\ |\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|_X\ \mathbf{Z}_q^{-(s_0\cdot\mathbf{A}+\mathbf{e}_0)}.$$

From Theorem 2 it follows that there exists a negligible function $\varepsilon(\lambda)$ such that

$$\mathcal{Z}_{\mathsf{LWE}^m_{n,q,\alpha q}}(|\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|_X) \ \approx_\varepsilon\ |\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|_X.$$

In other words, $\mathcal{A}$ receives as input a state in system X which is within negligible trace distance of the dual Gaussian state $|\hat{\psi}_{\mathbf{y}}\rangle$, which corresponds precisely to the input in $\mathsf{GaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(0)$.

If $(\mathbf{A}, \mathbf{b})$ is a uniformly random sample, where $\mathbf{b}$ is a random string $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$, then the adversary $\mathcal{A}$ receives as input the following quantum state in system X:

$$\mathcal{Z}(|\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|_X) = q^{-m}\sum_{\mathbf{u}\in\mathbb{Z}_q^m} \mathbf{Z}_q^{\mathbf{u}}\ |\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|_X\ \mathbf{Z}_q^{-\mathbf{u}}.$$

Because $\mathcal{Z}$ corresponds to the uniform Pauli-$\mathbf{Z}$ dephasing channel, it follows from Lemma 7 that

$$\mathcal{Z}(|\hat{\psi}_{\mathbf{y}}\rangle\langle\hat{\psi}_{\mathbf{y}}|_X) = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} |\langle\mathbf{x}|\hat{\psi}_{\mathbf{y}}\rangle|^2\ |\mathbf{x}\rangle\langle\mathbf{x}|_X.$$

In other words, $\mathcal{A}$ receives as input a mixed state which is the result of a computational basis measurement of the Gaussian state $|\hat{\psi}_{\mathbf{y}}\rangle$. Note that this corresponds precisely to the input in $\mathsf{GaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(1)$.

By assumption, the adversary $\mathcal{A}$ succeeds with advantage at least $1/p(\lambda)$. Therefore, the distinguisher $\mathcal{D}$ succeeds at the decisional $\mathsf{LWE}^m_{n,q,\alpha q}$ experiment with probability at least $1/p(\lambda) - \mathrm{negl}(\lambda)$. □

Theorem 5. *Let $n \in \mathbb{N}$ and $q \geq 2$ be a prime modulus with $m \geq 2n\log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\sigma \in (\sqrt{8m}, q/\sqrt{8m})$ be a function of λ and $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ be a matrix.*

Then, the following states are computationally indistinguishable assuming the quantum hardness of decisional $\mathsf{LWE}^m_{n,q,\alpha q}$, for any parameter $\alpha \in (0,1)$ with relative noise magnitude $\frac{q/\sigma}{\alpha q} = \lambda^{\omega(1)}$:

- *For any $(|\hat{\psi}_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenDual}(\mathbf{A}, \sigma)$ in Algorithm 1:*

$$|\hat{\psi}_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x}=\mathbf{y}\ (\mathrm{mod}\ q)}} \varrho_\sigma(\mathbf{x})\ |\mathbf{x}\rangle \quad \approx_c \quad |\mathbf{x}_0\rangle\ :\quad \mathbf{x}_0 \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\frac{\sigma}{\sqrt{2}}}.$$

- *For any $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenPrimal}(\mathbf{A}, \sigma)$ in Algorithm 2:*

$$|\psi_{\mathbf{y}}\rangle = \sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m} \varrho_{\frac{q}{\sigma}}(\mathbf{e})\,\omega_q^{-\langle\mathbf{s},\mathbf{y}\rangle}\,|\mathbf{s}\mathbf{A}+\mathbf{e}\rangle \quad \approx_c \quad \sum_{\mathbf{u}\in\mathbb{Z}_q^m} \omega_q^{-\langle\mathbf{u},\mathbf{x}_0\rangle}\,|\mathbf{u}\rangle\ :\quad \mathbf{x}_0 \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\frac{\sigma}{\sqrt{2}}}.$$

Moreover, the distribution of $\mathbf{y} \in \mathbb{Z}_q^n$ is negligibly close in total variation distance to the uniform distribution over $\mathbb{Z}_q^n$. Here, $\Lambda_q^{\mathbf{y}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\mathbf{x} = \mathbf{y} \pmod q\}$ denotes a coset of the lattice $\Lambda_q^\perp(\mathbf{A})$.

Proof. Let $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ be a random matrix. From Lemma 9 it follows that the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$ with overwhelming probability. Let us also recall the following simple facts about the discrete Gaussian. According to Lemma 10, the distribution of the syndrome $\mathbf{A}\cdot\mathbf{x} = \mathbf{y} \pmod q$ is statistically close to the uniform distribution over $\mathbb{Z}_q^n$, whenever $\mathbf{x} \sim D_{\mathbb{Z}^m,\sigma}$ and $\sigma = \omega(\sqrt{\log m})$. Moreover, the conditional distribution of $\mathbf{x} \sim D_{\mathbb{Z}^m,\sigma}$ given the syndrome $\mathbf{y} \in \mathbb{Z}_q^n$ is a discrete Gaussian distribution $D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}$.

Let us now show the first statement. Recall that in Theorem 4 we show that the Ajtaj hash function $h_{\mathbf{A}}(\mathbf{x}) = \mathbf{A}\cdot\mathbf{x} \pmod q$ is σ-Gaussian-collapsing assuming the decisional $\mathsf{LWE}^m_{n,q,\alpha q}$ assumption and a noise ratio $\frac{q/\sigma}{\alpha q} = \lambda^{\omega(1)}$. Therefore, for $\mathbf{y} \in \mathbb{Z}_q^n$, the (normalized variant of the) dual Gaussian state,

$$|\hat{\psi}_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x})\ |\mathbf{x}\rangle$$

is computationally indistinguishable from the (normalized) classical mixture,

$$\sum_{\mathbf{x}\in\mathbb{Z}_q^m} \left|\langle \mathbf{x}|\hat{\psi}_{\mathbf{y}}\rangle\right|^2 |\mathbf{x}\rangle\langle\mathbf{x}| = \left(\sum_{\substack{\mathbf{z}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{z}=\mathbf{y} \pmod q}} \varrho_{\sigma/\sqrt{2}}(\mathbf{z})\right)^{-1} \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_{\sigma/\sqrt{2}}(\mathbf{x})\ |\mathbf{x}\rangle\langle\mathbf{x}|\,,$$

which is the result of a computational basis measurement of $|\hat{\psi}_{\mathbf{y}}\rangle$.[4] Since $\sigma \in (\sqrt{8m}, q/\sqrt{8m})$, the tail bound in Lemma 11 implies that the above mixture is statistically close to the discrete Gaussian $D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\frac{\sigma}{\sqrt{2}}}$.

The second statement follows immediately by applying the (inverse) Fourier transform to both of the states above. Note that in Lemma 16 we showed that the primal Gaussian state

$$|\psi_{\mathbf{y}}\rangle = \sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m} \varrho_{\frac{q}{\sigma}}(\mathbf{e})\,\omega_q^{-\langle \mathbf{s},\mathbf{y}\rangle}\,|\mathbf{s}\mathbf{A}+\mathbf{e}\rangle$$

is within negligible trace distance of $\mathsf{FT}_q^\dagger|\hat{\psi}_{\mathbf{y}}\rangle$. This proves the claim.

□

5.2 Strong Gaussian-collapsing conjecture

Our quantum encryption schemes with certified deletion in Section 7 and Section 9 rely on the assumption that Ajtaj's hash function satisfies a strong Gaussian-collapsing property in the presence of leakage. We formalize the property as the following simple and falsifiable conjecture.

Conjecture (Strong Gaussian-Collapsing Conjecture)**.**
Let $\lambda \in \mathbb{N}$ be the security parameter, $n(\lambda) \in \mathbb{N}$, $q(\lambda) \geq 2$ be a modulus and $m \geq 2n\log q$ be an integer. Let $\sigma = \Omega(\sqrt{m})$ be a parameter and let $\mathcal{H} = \{H_\lambda\}_{\lambda\in\mathbb{N}}$ be the Ajtaj hash function family with

$$H_\lambda = \left\{h_{\mathbf{A}} : \mathbb{Z}_q^m \to \mathbb{Z}_q^n \ \textit{s.t.}\ h_{\mathbf{A}}(\mathbf{x}) = \mathbf{A}\cdot\mathbf{x} \pmod q;\ \mathbf{A}\in\mathbb{Z}_q^{n\times m}\right\}.$$

The Strong Gaussian-Collapsing Conjecture ($\mathsf{SGC}_{n,m,q,\sigma}$) states that, for every QPT adversary $\mathcal{A}$,

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

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

[4] Here, the additional factor $1/\sqrt{2}$ arises from the normalization of the dual Gaussian state $|\hat{\psi}_{\mathbf{y}}\rangle$.

1. *The challenger samples* $\bar{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{n\times(m-1)}$ *and prepares the quantum state*

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

where $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}}\cdot\bar{\mathbf{x}} \pmod q] \in \mathbb{Z}_q^{n\times m}$ *is a matrix with* $\bar{\mathbf{x}} \xleftarrow{\$} \{0,1\}^{m-1}$.

2. *The challenger measures system* Y *in the computational basis, resulting in the state*

$$|\hat{\psi}_{\mathbf{y}}\rangle_{XY} = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle_X \otimes |\mathbf{y}\rangle_Y \,.$$

3. *If* $b = 0$, *the challenger does nothing. Else, if* $b = 1$, *the challenger measures system* X *of the quantum state* $|\hat{\psi}_{\mathbf{y}}\rangle$ *in the computational basis. Finally, the challenger sends the outcome state in systems* X *to* $\mathcal{A}$, *together with the matrix* $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ *and the string* $\mathbf{y} \in \mathbb{Z}_q^n$.

4. $\mathcal{A}$ *sends a classical witness* $\mathbf{w} \in \mathbb{Z}_q^m$ *to the challenger.*

5. *The challenger checks whether* $\mathbf{A}\cdot\mathbf{w} = \mathbf{y} \pmod q$ *and* $\|\mathbf{w}\| \leq \sqrt{m}\sigma/\sqrt{2}$. *If* $\mathbf{w}$ *passes both checks, the challenger sends* $\mathbf{t} = (\bar{\mathbf{x}}, -1) \in \mathbb{Z}_q^m$ *to* $\mathcal{A}$ *with* $\mathbf{A}\cdot\mathbf{t} = \mathbf{0} \pmod q$. *Else, the challenger aborts.*

6. $\mathcal{A}$ *returns a bit* b', *which we define as the output of the experiment.*

Remark. *We also consider an* N*-fold variant of* $\mathsf{SGC}_{n,m,q,\sigma}$, *which we denote by* $\mathsf{SGC}^N_{n,m,q,\sigma}$, *where the challenger prepares* N *independent states* $|\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes |\hat{\psi}_{\mathbf{y}_N}\rangle$ *in Steps* 1–2, *for outcomes* $\mathbf{y}_1, \ldots, \mathbf{y}_N \in \mathbb{Z}_q^n$. *A simple hybrid argument shows that* $\mathsf{SGC}^N_{n,m,q,\sigma}$ *is implied by* $\mathsf{SGC}_{n,m,q,\sigma}$, *for any* $N = \mathrm{poly}(\lambda)$.

Towards a proof of the strong-Gaussian-collapsing conjecture. Unfortunately, we currently do not know how to prove Conjecture 5.2 from standard assumptions, such as LWE or ISIS. The difficulty emerges when we attempt to reduce the security to the LWE (or ISIS) problem with respect to the same matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$. In order to simulate the experiment $\mathsf{StrongGaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}$ with respect to an adversary $\mathcal{A}$, we have to eventually forward a short trapdoor vector $\mathbf{t} \in \mathbb{Z}^m$ in order to simulate the second phase of the experiment once $\mathcal{A}$ has produced a valid witness. Notice, however, that the reduction has no way of obtaining a short vector $\mathbf{t}$ in the kernel of $\mathbf{A}$ as it is trying to break the underlying LWE (or ISIS) problem with respect to $\mathbf{A}$ in the first place. Therefore, any successful security proof must necessarily exploit the fact that there is *interaction* between the challenger and the adversary $\mathcal{A}$, and that a short trapdoor vector $\mathbf{t}$ is only revealed *after* $\mathcal{A}$ has already produced a valid short pre-image of $\mathbf{y} \in \mathbb{Z}_q^n$.

When trying to distinguish between the state $|\hat{\psi}_{\mathbf{y}}\rangle$ and a single Gaussian pre-image $|\mathbf{x}_0\rangle$ with the property that $\mathbf{A}\cdot\mathbf{x}_0 = \mathbf{y} \pmod q$, it is useful to work with the Fourier basis. Without loss of generality, we can assume that $\mathcal{A}$ instead receives one of the following states during in Step 2; namely

$$\sum_{\mathbf{s}\in\mathbb{Z}_q^n} \sum_{\mathbf{e}\in\mathbb{Z}_q^m} \varrho_{\frac{q}{\sigma}}(\mathbf{e}) \, \omega_q^{-\langle \mathbf{s},\mathbf{y}\rangle} \, |\mathbf{s}\mathbf{A}+\mathbf{e}\rangle_X \qquad \text{or} \qquad \sum_{\mathbf{u}\in\mathbb{Z}_q^m} \omega_q^{-\langle \mathbf{u},\mathbf{x}_0\rangle} \, |\mathbf{u}\rangle_X \,.$$

One natural approach is prepare an auxiliary system, say B, which could later help the adversary determine whether X corresponds to a superposition of LWE samples or a superposition of uniform samples once the trapdoor $\mathbf{t}$ is revealed (ideally, without disturbing X so as to allow for a Fourier basis measurement).

Because finding a valid witness $\mathbf{w}$ to the ISIS problem specified by $(\mathbf{A}, \mathbf{y})$ now amounts to a Fourier basis projection (as in Definition 17), the entropic uncertainty relation in Theorem 3 immediately rules out large class of attacks, including the *shift-by-*LWE*-sample attack* we described in Section 1.2. There, the idea is to reversibly shift system X by a fresh LWE sample into an auxiliary system B. If system X corresponds to a superposition of LWE samples, we obtain a separate LWE sample which is *re-randomized*, whereas, if X is a superposition of uniform samples, the outcome remains random. Hence, if the aforementioned procedure succeeded without disturbing system X, we could potentially find a valid witness $\mathbf{w}$ and simultaneously distinguish the auxiliary system B with access to the trapdoor $\mathbf{t}$. As we observed before, however, such an attack must necessarily entangle the two systems X and B in a way that prevents it from finding a solution to the ISIS problem specified by $(\mathbf{A}, \mathbf{y})$. Intuitively, if the state in system X yields a short-pre image $\mathbf{w}$ *with high probability* via a Fourier basis measurement, then system X cannot be entangled with any auxiliary systems. Because the set $\mathcal{S}$ of valid short pre-images (i.e. the set of solution to the ISIS problem specified by $\mathbf{A}$ and $\mathbf{y}$) is much smaller than the size of $\mathbb{Z}_q^m$ (in particular, if $\sigma\sqrt{m} \ll q$), Theorem 3 tells us that the min-entropy of system X (once it is measured in the computational basis) given system B must necessarily be large. We remark that this statement holds *information-theoretically*, and does not rely on the hardness of LWE. This suggests that, even if the trapdoor $\mathbf{t}$ is later revealed, system B cannot contain any relevant information about whether system X initially corresponded to a superposition of LWE samples, or to a superposition of uniform samples. While this argument is not sufficient to prove Conjecture 5.2, it captures the inherent difficulty in extracting information encoded in two mutually unbiased bases, i.e. the computational basis and the Fourier basis.

6 Public-Key Encryption with Certified Deletion

In this section, we formalize the notion of public-key encryption with certified deletion.

6.1 Definition

We first introduce the following definition.

Definition 20 (Public-key encryption with certified deletion)**.** *A public-key encryption scheme with certified deletion (*$\mathsf{PKE}_{\mathsf{CD}}$*)* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$ *with plaintext space* $\mathcal{M}$ *consists of a tuple of* QPT *algorithms, a key generation algorithm* KeyGen*, an encryption algorithm* Enc*, and a decryption algorithm* Dec*, a deletion algorithm* Del*, and a verification algorithm* Vrfy*.*

$\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: *takes as input the parameter* 1^λ *and outputs a public key* pk *and secret key* sk*.*

$\mathsf{Enc}(\mathsf{pk}, m) \to (\mathsf{vk}, \mathsf{CT})$: *takes as input the public key* pk *and a plaintext* $m \in \mathcal{M}$*, and outputs a classical verification key* vk *together with a quantum ciphertext* CT*.*

$\mathsf{Dec}(\mathsf{sk}, \mathsf{CT}) \to m'$ **or** $\bot$: *takes as input the secret key* sk *and ciphertext* CT*, and outputs* $m' \in \mathcal{M}$ *or* $\bot$*.*

$\mathsf{Del}(\mathsf{CT}) \to \pi$: *takes as input a ciphertext* CT *and outputs a classical certificate* π*.*

$\mathsf{Vrfy}(\mathsf{vk}, \pi) \to \top$ **or** $\bot$: *takes as input the verification key* vk *and certificate* π*, and outputs* $\top$ *or* $\bot$*.*

Definition 21 (Correctness of $\mathsf{PKE}_{\mathsf{CD}}$)**.** *We require two separate kinds of correctness properties, one for decryption and one for verification.*

(Decryption correctness:) For any $\lambda \in \mathbb{N}$, and for any $m \in \mathcal{M}$:

$$\Pr\left[\mathsf{Dec}(\mathsf{sk},\mathsf{CT}) \neq m \;\middle|\; \begin{smallmatrix}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KeyGen}(1^\lambda)\\ \mathsf{CT}\leftarrow\mathsf{Enc}(\mathsf{pk},m)\end{smallmatrix}\right] \leq \mathrm{negl}(\lambda).$$

(Verification correctness:) For any $\lambda \in \mathbb{N}$, and for any $m \in \mathcal{M}$:

$$\Pr\left[\mathsf{Vrfy}(\mathsf{vk},\pi) = \bot \;\middle|\; \begin{smallmatrix}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KeyGen}(1^\lambda)\\ (\mathsf{vk},\mathsf{CT})\leftarrow\mathsf{Enc}(\mathsf{pk},m)\\ \pi\leftarrow\mathsf{Del}(\mathsf{CT})\end{smallmatrix}\right] \leq \mathrm{negl}(\lambda).$$

The notion of IND-CPA-CD security for public-key encryption was first introduced by Hiroka, Morimae, Nishimaki and Yamakawa [HMNY21b].

6.2 Certified deletion security

In terms of security, we adopt the following definition.

Definition 22 (Certified deletion security for PKE)**.** *Let $\Sigma = (\mathsf{KeyGen},\mathsf{Enc},\mathsf{Dec},\mathsf{Del},\mathsf{Vrfy})$ be a $\mathsf{PKE}_{\mathsf{CD}}$ scheme and let $\mathcal{A}$ be a QPT adversary (in terms of the security parameter $\lambda \in \mathbb{N}$). We define the security experiment $\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(b)$ between $\mathcal{A}$ and a challenger as follows:*

1. *The challenger generates a pair $(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$, and sends pk to $\mathcal{A}$.*
2. *$\mathcal{A}$ sends a plaintext pair $(m_0,m_1) \in \mathcal{M}\times\mathcal{M}$ to the challenger.*
3. *The challenger computes $(\mathsf{vk},\mathsf{CT}_b) \leftarrow \mathsf{Enc}(\mathsf{pk},m_b)$, and sends CT_b to $\mathcal{A}$.*
4. *At some point in time, $\mathcal{A}$ sends the certificate π to the challenger.*
5. *The challenger computes $\mathsf{Vrfy}(\mathsf{vk},\pi)$ and sends sk to $\mathcal{A}$, if the output is $\top$, and sends $\bot$ otherwise.*
6. *$\mathcal{A}$ outputs a guess $b' \in \{0,1\}$, which is also the output of the experiment.*

We say that the scheme Σ is IND-CPA-CD-secure if, for any QPT adversary $\mathcal{A}$, it holds that

$$\mathsf{Adv}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(0)=1] - \Pr[\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1)=1]| \leq \mathrm{negl}(\lambda).$$

7 Dual-Regev Public-Key Encryption with Certified Deletion

In this section, we consider the Dual-Regev PKE scheme due to Gentry, Peikert and Vaikuntanathan [GPV07]. Unlike Regev's original PKE scheme in [Reg05], the Dual-Regev PKE scheme has the useful property that the ciphertext takes the form of a regular sample from the LWE distribution together with an additive shift which depends on the plaintext.

7.1 Construction

Parameters. Let $\lambda \in \mathbb{N}$ be the security parameter. We choose the following set of parameters for our Dual-Regev PKE scheme with certified deletion (each parameterized by λ).

- an integer $n \in \mathbb{N}$.
- a prime modulus $q \geq 2$.
- an integer $m \geq 2n \log q$.
- a noise ratio $\alpha \in (0,1)$ such that $\sqrt{8(m+1)} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)}}$.

Construction 1 (Dual-Regev PKE with Certified Deletion)**.** *Let $\lambda \in \mathbb{N}$. The Dual-Regev* PKE *scheme* $\mathsf{DualPKE_{CD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$ *with certified deletion is defined as follows:*

$\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: *sample a random matrix* $\bar{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ *and a vector* $\bar{\mathbf{x}} \xleftarrow{\$} \{0,1\}^m$ *and choose* $\mathbf{A} = [\bar{\mathbf{A}} | \bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q]$. *Output* $(\mathsf{pk}, \mathsf{sk})$, *where* $\mathsf{pk} = \mathbf{A} \in \mathbb{Z}_q^{n \times (m+1)}$ *and* $\mathsf{sk} = (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$.

$\mathsf{Enc}(\mathsf{pk}, x) \to (\mathsf{vk}, |\mathsf{CT}\rangle)$*: parse* $\mathbf{A} \leftarrow \mathsf{pk}$ *and run* $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenPrimal}(\mathbf{A}, 1/\alpha)$ *in Algorithm 2, where* $\mathbf{y} \in \mathbb{Z}_q^n$. *To encrypt a single bit* $b \in \{0,1\}$, *output the pair*

$$\left(\mathsf{vk} \leftarrow (\mathbf{A} \in \mathbb{Z}_q^{n \times (m+1)}, \mathbf{y} \in \mathbb{Z}_q^n), \quad |\mathsf{CT}\rangle \leftarrow \mathbf{X}_q^{(0,\ldots,0,b \cdot \lfloor \frac{q}{2} \rfloor)} |\psi_{\mathbf{y}}\rangle\right),$$

where vk *is the public verification key and* $|\mathsf{CT}\rangle$ *is an* $(m+1)$*-qudit quantum ciphertext.*

$\mathsf{Dec}(\mathsf{sk}, |\mathsf{CT}\rangle) \to \{0,1\}$: *to decrypt, measure the ciphertext* $|\mathsf{CT}\rangle$ *in the computational basis with outcome* $\mathbf{c} \in \mathbb{Z}_q^m$. *Compute* $\mathsf{sk}^T \cdot \mathbf{c} \in \mathbb{Z}_q$ *and output* 0, *if it is closer to* 0 *than to* $\lfloor \frac{q}{2} \rfloor$, *and output* 1, *otherwise.*

$\mathsf{Del}(|\mathsf{CT}\rangle) \to \pi$: *Measure* $|\mathsf{CT}\rangle$ *in the Fourier basis and output the measurement outcome* $\pi \in \mathbb{Z}_q^{m+1}$.

$\mathsf{Vrfy}(\mathsf{vk}, \pi) \to \{\top, \bot\}$: *to verify a deletion certificate* $\pi \in \mathbb{Z}_q^{m+1}$, *parse* $(\mathbf{A}, \mathbf{y}) \leftarrow \mathsf{vk}$ *and output* $\top$, *if* $\mathbf{A} \cdot \pi = \mathbf{y} \pmod q$ *and* $\|\pi\| \leq \sqrt{m+1}/\sqrt{2}\alpha$, *and output* $\bot$, *otherwise.*

Proof of correctness. Let us now establish the correctness properties of $\mathsf{DualPKE_{CD}}$ in Construction 1.

Lemma 17 (Correctness of decryption)**.** *Let $n \in \mathbb{N}$ and $q \geq 2$ be a prime modulus with $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let α be a ratio with $\sqrt{8(m+1)} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)}}$. Then, for $b \in \{0,1\}$, the scheme* $\mathsf{DualPKE_{CD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$ *in Construction 1 satisfies:*

$$\Pr\left[\mathsf{Dec}(\mathsf{sk}, |\mathsf{CT}\rangle) = b \;\middle|\; \begin{matrix}(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ (\mathsf{vk}, |\mathsf{CT}\rangle) \leftarrow \mathsf{Enc}(\mathsf{pk}, b)\end{matrix}\right] \geq 1 - \mathrm{negl}(\lambda).$$

Proof. By the Leftover Hash Lemma (Lemma 4), the distribution of $\mathbf{A} = [\bar{\mathbf{A}} | \bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q]$ is within negligible total variation distance of the uniform distribution over $\mathbb{Z}_q^{n \times (m+1)}$. Moreover, from Lemma 9 it follows that the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$ with overwhelming probability. Since the noise ratio $\alpha \in (0,1)$

satisfies $\sqrt{8(m+1)} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)}}$, it then follows from Corollary 1 that the ciphertext $|\mathsf{CT}\rangle$ is within negligible trace distance of the state

$$\sum_{\mathbf{s}\in\mathbb{Z}_q^n} \sum_{\mathbf{e}\in\mathbb{Z}_q^{m+1}} \varrho_{\alpha q}(\mathbf{e})\, \omega_q^{-\langle \mathbf{s},\mathbf{y}\rangle}\, |\mathbf{s}\mathbf{A} + \mathbf{e} + (0,\ldots,0,b\cdot\lfloor\frac{q}{2}\rfloor)\rangle$$

A measurement in computational basis yields an outcome $\mathbf{c}$ such that

$$\mathbf{c} = \mathbf{s}_0\mathbf{A} + \mathbf{e}_0 + (0,\ldots,0,b\cdot\lfloor\frac{q}{2}\rfloor) \in \mathbb{Z}_q^{m+1},$$

where $\mathbf{s}_0 \xleftarrow{\$} \mathbb{Z}_q^n$ is random and where $\mathbf{e}_0 \sim D_{\mathbb{Z}_q^{m+1},\frac{\alpha q}{\sqrt{2}}}$ is a sample from the (truncated) discrete Gaussian such that $\|\mathbf{e}_0\| \leq \alpha q\sqrt{\frac{m+1}{2}} < \lfloor\frac{q}{4}\rfloor$. Since $\mathsf{Dec}(\mathsf{sk},|\mathsf{CT}\rangle)$ computes $\mathsf{sk}^T\cdot\mathbf{c} \in \mathbb{Z}\cap(-\frac{q}{2},\frac{q}{2}]$ and outputs 0, if it is closer to 0 than to $\lfloor\frac{q}{2}\rfloor$ over , and 1 otherwise, it succeeds with overwhelming probability.

□

Let us now prove the following property.

Lemma 18 (Correctness of verification)**.** *Let $n \in \mathbb{N}$ and $q \geq 2$ be a prime modulus with $m \geq 2n\log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let α be a ratio with $\sqrt{8(m+1)} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)}}$. Then, for $b \in \{0,1\}$, the scheme $\mathsf{DualPKE}_{\mathsf{CD}} = (\mathsf{KeyGen},\mathsf{Enc},\mathsf{Dec},\mathsf{Del},\mathsf{Vrfy})$ in Construction 1 satisfies:*

$$\Pr\left[\mathsf{Verify}(\mathsf{vk},\pi) = \top \;\middle|\; \begin{smallmatrix}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KeyGen}(1^\lambda)\\ (\mathsf{vk},|\mathsf{CT}\rangle)\leftarrow\mathsf{Enc}(\mathsf{pk},b)\\ \pi\leftarrow\mathsf{Del}(|\mathsf{CT}\rangle)\end{smallmatrix}\right] \geq 1-\mathrm{negl}(\lambda).$$

Proof. By the Leftover Hash Lemma (Lemma 4), the distribution of $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}}\cdot\bar{\mathbf{x}} \pmod q]$ is within negligible total variation distance of the uniform distribution over $\mathbb{Z}_q^{n\times(m+1)}$. From Lemma 9 it follows that the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$ with overwhelming probability. Since $\alpha \in (0,1)$ is a ratio parameter with $\sqrt{8(m+1)} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)}}$, Corollary 1 implies that the Fourier transform of the ciphertext $|\mathsf{CT}\rangle$ is within negligible trace distance of the state

$$|\widehat{\mathsf{CT}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^{m+1}:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_{1/\alpha}(\mathbf{x})\, \omega_q^{\langle \mathbf{x},(0,\ldots,0,b\cdot\lfloor\frac{q}{2}\rfloor)\rangle}\, |\mathbf{x}\rangle\,.$$

From Lemma 11, it follows that the distribution of computational basis measurement outcomes is within negligible total variation distance of $\pi \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\frac{1}{\sqrt{2}\alpha}}$ with $\|\pi\| \leq \sqrt{m+1}/\sqrt{2}\alpha$. This proves the claim.

□

7.2 Proof of security

Let us now analyze the security of our Dual-Regev PKE scheme with certified deletion in Construction 1.

IND-CPA security of $\mathsf{DualPKE_{CD}}$. We first prove that our public-key encryption scheme $\mathsf{DualPKE_{CD}}$ in Construction 1 satisfies the notion IND-CPA security according to Definition 12. The proof follows from Theorem 5 and assumes the hardness of (decisional) LWE (Definition 15). We add it for completeness.

Theorem 6. *Let $n \in \mathbb{N}$ and $q \geq 2$ be a prime modulus with $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\alpha \in (0,1)$ be a noise ratio parameter with $\sqrt{8(m+1)} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)}}$. Then, the scheme $\mathsf{DualPKE_{CD}}$ in Construction 1 is IND-CPA-secure assuming the quantum hardness of the decisional $\mathsf{LWE}^m_{n,q,\beta q}$ problem, for any $\beta \in (0,1)$ with $\alpha/\beta = \lambda^{\omega(1)}$.*

Proof. Let $\Sigma = \mathsf{DualPKE_{CD}}$. We need to show that, for any QPT adversary $\mathcal{A}$, it holds that

$$\mathsf{Adv}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

Consider the experiment $\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(b)$ between the adversary $\mathcal{A}$ and a challenger taking place as follows:

1. The challenger generates a pair $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$, and sends pk to $\mathcal{A}$.
2. $\mathcal{A}$ sends a distinct plaintext pair $(m_0, m_1) \in \{0,1\} \times \{0,1\}$ to the challenger.
3. The challenger computes $(\mathsf{vk}, \mathsf{CT}_b) \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b)$, and sends $|\mathsf{CT}_b\rangle$ to $\mathcal{A}$.
4. $\mathcal{A}$ outputs a guess $b' \in \{0,1\}$, which is also the output of the experiment.

Recall that the procedure $\mathsf{Enc}(\mathsf{pk}, m_b)$ outputs a pair $(\mathsf{vk}, |\mathsf{CT}_b\rangle)$, where $(\mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{vk}$ is the verification key and where the ciphertext $|\mathsf{CT}_b\rangle$ is within negligible trace distance of

$$\sum_{\mathbf{s}\in\mathbb{Z}_q^n} \sum_{\mathbf{e}\in\mathbb{Z}_q^{m+1}} \varrho_{\alpha q}(\mathbf{e})\, \omega_q^{-\langle \mathbf{s},\mathbf{y}\rangle}\, |\mathbf{sA} + \mathbf{e} + (0,\ldots,0, m_b \cdot \lfloor q/2 \rfloor) \pmod q\rangle \tag{16}$$

Let $\beta \in (0,1)$ be such that $\alpha/\beta = \lambda^{\omega(1)}$. From Theorem 5 it follows that, under the (decisional) $\mathsf{LWE}^m_{n,q,\beta q}$ assumption, the quantum ciphertext $|\mathsf{CT}_b\rangle$ is computationally indistinguishable from the state

$$\sum_{\mathbf{u}\in\mathbb{Z}_q^{m+1}} \omega_q^{-\langle \mathbf{u},\mathbf{x}_0\rangle}\, |\mathbf{u}\rangle\,, \qquad \mathbf{x}_0 \sim D_{\Lambda^{\mathbf{y}}_q(\mathbf{A}),\frac{1}{\sqrt{2}\alpha}}. \tag{17}$$

Because the state in Eq. (20) is completely independent of $b \in \{0,1\}$, it follows that

$$\mathsf{Adv}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

This proves the claim. □

IND-CPA-CD security of $\mathsf{DualPKE_{CD}}$. In this section, we prove that our public-key encryption scheme $\mathsf{DualPKE_{CD}}$ in Construction 1 satisfies the notion of *certified deletion security* assuming the *Strong Gaussian-Collapsing (SGC) Conjecture* (see Conjecture 5.2). This is a strengthening of the Gaussian-collapsing property which we proved under the (decisional) LWE assumption (see Theorem 4).

Theorem 7. *Let $n \in \mathbb{N}$ and $q \geq 2$ be a prime modulus with $m \geq 2n \log q$, each parameterized by $\lambda \in \mathbb{N}$. Let Let α be a ratio with $\sqrt{8(m+1)} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)}}$. Then, the scheme $\mathsf{DualPKE_{CD}}$ in Construction 1 is IND-CPA-CD-secure assuming the Strong Gaussian-Collapsing property $\mathsf{SGC}_{n,m+1,q,\frac{1}{\alpha}}$ from Conjecture 5.2.*

Proof. Let $\Sigma = \mathsf{DualPKE}_{\mathsf{CD}}$. We need to show that, for any QPT adversary $\mathcal{A}$, it holds that

$$\mathsf{Adv}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

We consider the following sequence of hybrids:

$\mathbf{H_0}$: This is the experiment $\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(0)$ between $\mathcal{A}$ and a challenger:

1. The challenger samples a random matrix $\bar{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and a vector $\bar{\mathbf{x}} \xleftarrow{\$} \{0,1\}^m$ and chooses $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}}\cdot\bar{\mathbf{x}} \pmod q]$. The challenger chooses the secret key $\mathsf{sk} \leftarrow (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$ and the public key $\mathsf{pk} \leftarrow \mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}$.
2. $\mathcal{A}$ sends a distinct plaintext pair $(m_0, m_1) \in \{0,1\} \times \{0,1\}$ to the challenger. (Note: Without loss of generality, we can just assume that $m_0 = 0$ and $m_1 = 1$).
3. The challenger runs $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenPrimal}(\mathbf{A}, 1/\alpha)$ in Algorithm 2, and outputs
$$\left(\mathsf{vk} \leftarrow (\mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}, \mathbf{y} \in \mathbb{Z}_q^n), \quad |\mathsf{CT}_0\rangle \leftarrow |\psi_{\mathbf{y}}\rangle\right).$$
4. At some point in time, $\mathcal{A}$ returns a certificate π to the challenger.
5. The challenger verifies π and outputs $\top$, if $\mathbf{A}\cdot\pi = \mathbf{y} \pmod q$ and $\|\pi\| \leq \sqrt{m+1}/\sqrt{2}\alpha$, and output $\bot$, otherwise. If π passes the test with outcome $\top$, the challenger sends sk to $\mathcal{A}$.
6. $\mathcal{A}$ outputs a guess $b' \in \{0,1\}$, which is also the output of the experiment.

$\mathbf{H_1}$: This is same experiment as in $\mathbf{H_0}$, except that (in Step 3) the challenger prepares the ciphertext in the Fourier basis rather than the standard basis. In other words, $\mathcal{A}$ receives the pair

$$\left(\mathsf{vk} \leftarrow (\mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}, \mathbf{y} \in \mathbb{Z}_q^n), \quad |\mathsf{CT}_0\rangle \leftarrow \mathsf{FT}_q|\psi_{\mathbf{y}}\rangle\right).$$

$\mathbf{H_2}$: This is the experiment $\mathsf{StrongGaussCollapseExp}_{\mathcal{H},\mathcal{D},\lambda}(0)$ in Conjecture 5.2:

1. The challenger samples a random matrix $\bar{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and a vector $\bar{\mathbf{x}} \xleftarrow{\$} \{0,1\}^m$ and chooses $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}}\cdot\bar{\mathbf{x}} \pmod q]$ and $\mathbf{t} = (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$.
2. The challenger runs $(|\hat{\psi}_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenDual}(\mathbf{A}, \sigma)$ in Algorithm 1, where $\mathbf{y} \in \mathbb{Z}_q^n$, and sends the triplet $(|\hat{\psi}_{\mathbf{y}}\rangle, \mathbf{A}, \mathbf{y})$ to the adversary $\mathcal{A}$.
3. At some point in time, $\mathcal{A}$ returns a certificate π to the challenger.
4. The challenger verifies π and outputs $\top$, if $\mathbf{A}\cdot\pi = \mathbf{y} \pmod q$ and $\|\pi\| \leq \sqrt{m+1}/\sqrt{2}\alpha$, and output $\bot$, otherwise. If π passes the test with outcome $\top$, the challenger sends $\mathbf{t}$ to $\mathcal{A}$.
5. $\mathcal{A}$ outputs a guess $b' \in \{0,1\}$, which is also the output of the experiment.

$\mathbf{H_3}$: This is the experiment $\mathsf{StrongGaussCollapseExp}_{\mathcal{H},\mathcal{D},\lambda}(1)$ in Conjecture 5.2; it is the same as $\mathbf{H_2}$, except that the state $|\hat{\psi}_{\mathbf{y}}\rangle$ (in Step 2) is measured in the computational basis before it is sent to $\mathcal{A}$.

$\mathbf{H_4}$: This is same experiment as $\mathbf{H_3}$, except that (in Step 2) the challenger additionally applies the Pauli operator $\mathbf{Z}_q^{(0,\ldots,0,\lfloor\frac{q}{2}\rfloor)}$ to the state $|\hat{\psi}_{\mathbf{y}}\rangle$ before it is measured in the computational basis.

$\mathbf{H_5}$: This is same experiment as $\mathbf{H_4}$, except that (in Step 2) $\mathcal{A}$ receives the triplet

$$(\mathbf{Z}_q^{(0,\ldots,0,\lfloor \frac{q}{2} \rfloor)} |\hat{\psi}_{\mathbf{y}}\rangle, \quad \mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}, \quad \mathbf{y} \in \mathbb{Z}_q^n).$$

$\mathbf{H_6}$: This is same experiment as $\mathbf{H_5}$, except that (in Step 2) the challenger prepares the quantum state $\mathbf{Z}_q^{(0,\ldots,0,\lfloor \frac{q}{2} \rfloor)} |\hat{\psi}_{\mathbf{y}}\rangle$ in the (inverse) Fourier basis instead. In other words, $\mathcal{A}$ receives the triplet

$$(\mathsf{FT}_q^\dagger \mathbf{Z}_q^{(0,\ldots,0,\lfloor \frac{q}{2} \rfloor)} |\hat{\psi}_{\mathbf{y}}\rangle, \quad \mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}, \quad \mathbf{y} \in \mathbb{Z}_q^n).$$

$\mathbf{H_7}$: This is the experiment $\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1)$.

We now show that the hybrids are indistinguishable.

Claim 2.

$$\Pr[\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(0) = 1] = \Pr[\mathbf{H_1} = 1].$$

Proof. Without loss of generality, we can assume that $\mathcal{A}$ applies the inverse Fourier transform immediately upon receiving the quantum ciphertext. Therefore, the success probabilities are identical in $\mathbf{H_0}$ and $\mathbf{H_1}$. □

Claim 3.

$$\Pr[\mathbf{H_1} = 1] = \Pr[\mathbf{H_2} = 1].$$

Proof. Because the challenger in $\mathbf{H_1}$ always sends the ciphertext $|\mathsf{CT}_0\rangle$ corresponding to $m_0 = 0$ to the adversary $\mathcal{A}$, the two hybrids $\mathbf{H_1}$ and $\mathbf{H_2}$ are identical. □

Claim 4. *Under the Strong Gaussian-Collapsing property* $\mathsf{SGC}_{n,m+1,q,\frac{1}{\alpha}}$*, it holds that*

$$|\Pr[\mathbf{H_2} = 1] - \Pr[\mathbf{H_3} = 1]| \leq \mathrm{negl}(\lambda).$$

Proof. This follows directly from Conjecture 5.2. □

Claim 5.

$$\Pr[\mathbf{H_3} = 1] = \Pr[\mathbf{H_4} = 1].$$

Proof. Because the challenger measures the state $|\hat{\psi}_{\mathbf{y}}\rangle$ in Step 2 in the computational basis, applying the phase operator $\mathbf{Z}_q^{(0,\ldots,0,\lfloor \frac{q}{2} \rfloor)}$ before the measurement does not affect the measurement outcome. □

Claim 6. *Under the Strong Gaussian-Collapsing property* $\mathsf{SGC}_{n,m+1,q,\frac{1}{\alpha}}$*, it holds that*

$$|\Pr[\mathbf{H_4} = 1] - \Pr[\mathbf{H_5} = 1]| \leq \mathrm{negl}(\lambda).$$

Proof. This follows from Conjecture 5.2 since, without loss of generality, we can assume that $\mathcal{A}$ applies the phase operator $\mathbf{Z}_q^{(0,\ldots,0,\lfloor \frac{q}{2} \rfloor)}$ immediately upon receiving the state $|\hat{\psi}_{\mathbf{y}}\rangle$. □

Claim 7.

$$\Pr[\mathbf{H_5} = 1] = \Pr[\mathbf{H_6} = 1].$$

Proof. Without loss of generality, we can assume that $\mathcal{A}$ applies the Fourier transform immediately upon receiving the state $\mathbf{Z}_q^{(0,\ldots,0,\lfloor \frac{q}{2} \rfloor)} |\hat{\psi}_{\mathbf{y}}\rangle$. Therefore, the success probabilities are identical in $\mathbf{H_5}$ and $\mathbf{H_6}$. □

Claim 8.

$$|\Pr[\mathbf{H_6}=1]-\Pr[\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1)=1]|\leq \mathrm{negl}(\lambda).$$

Proof. From Lemma 6, we have $\mathsf{FT}_q\mathbf{X}^{\mathbf{v}}_q=\mathbf{Z}^{\mathbf{v}}_q\mathsf{FT}_q$, for all $\mathbf{v}\in\mathbb{Z}^m_q$. Hence, in $\mathbf{H_6}$, we can instead assume that the challenger runs $(|\psi_{\mathbf{y}}\rangle,\mathbf{y})\leftarrow\mathsf{GenPrimal}(\mathbf{A},1/\alpha)$ in Algorithm 2 and sends the following to $\mathcal{A}$:

$$\left(\mathsf{vk}\leftarrow(\mathbf{A}\in\mathbb{Z}^{n\times(m+1)}_q,\mathbf{y}\in\mathbb{Z}^n_q),\quad |\mathsf{CT}_1\rangle\leftarrow\mathbf{X}^{(0,\ldots,0,\lfloor\frac{q}{2}\rfloor)}_q|\psi_{\mathbf{y}}\rangle\right).$$

From Corollary 1, we have that $\mathsf{FT}^{\dagger}_q\mathbf{Z}^{\mathbf{v}}_q|\hat{\psi}_{\mathbf{y}}\rangle$ and $\mathbf{X}^{\mathbf{v}}_q|\psi_{\mathbf{y}}\rangle$ are within negligible trace distance, for all $\mathbf{v}\in\mathbb{Z}^m_q$. Because the challenger in $\mathbf{H_7}$ always sends the ciphertext $|\mathsf{CT}_1\rangle$ corresponding to $m_1=1$ to the adversary $\mathcal{A}$, it follows that the distinguishing advantage between $\mathbf{H_6}$ and $\mathbf{H_7}=\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1)$ is negligible. □

Because the hybrids $\mathbf{H_0}$ and $\mathbf{H_7}$ are indistinguishable, this implies that

$$\mathsf{Adv}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A}}(\lambda)\leq\mathrm{negl}(\lambda).$$

□

Next, we show how to extend our Dual-Regev PKE scheme with certified deletion in Construction 1 to a fully homomorphic encryption scheme of the same type.

8 Fully Homomorphic Encryption with Certified Deletion

In this section, we formalize the notion of homomorphic encryption with certified deletion which enables an untrusted quantum server to compute on encrypted data and, if requested, to simultaneously prove data deletion to a client. We also provide several notions of certified deletion security.

8.1 Definition

We begin with the following definition.

Definition 23 (Homomorphic encryption with certified deletion)**.** *A homomorphic encryption scheme with certified deletion is a tuple* $\mathsf{HE}_{\mathsf{CD}}=(\mathsf{KeyGen},\mathsf{Enc},\mathsf{Dec},\mathsf{Eval},\mathsf{Del},\mathsf{Vrfy})$ *of* QPT *algorithms (in the security parameter* $\lambda\in\mathbb{N}$*), a key generation algorithm* KeyGen*, an encryption algorithm* Enc*, a decryption algorithm* Dec*, an evaluation algorithm* Eval*, a deletion algorithm* Del*, and a verification algorithm* Vrfy*.*

$\mathsf{KeyGen}(1^\lambda)\rightarrow(\mathsf{pk},\mathsf{sk})$: *takes as input* 1^λ *and outputs a public key* pk *and secret key* sk*.*

$\mathsf{Enc}(\mathsf{pk},x)\rightarrow(\mathsf{vk},\mathsf{CT})$: *takes as input the public key* pk *and a plaintext* $x\in\{0,1\}$*, and outputs a classical verification key* vk *together with a quantum ciphertext* CT*.*

$\mathsf{Dec}(\mathsf{sk},\mathsf{CT})\rightarrow x'$ **or** $\bot$: *takes as input a key* sk *and ciphertext* CT*, and outputs* $x'\in\{0,1\}$ *or* $\bot$*.*

$\mathsf{Eval}(C,\mathsf{CT},\mathsf{pk})\rightarrow\widetilde{\mathsf{CT}}$*: takes as input a key* pk *and applies a circuit* $C:\{0,1\}^\ell\rightarrow\{0,1\}$ *to a product of quantum ciphertexts* $\mathsf{CT}=\mathsf{CT}_1\otimes\cdots\otimes\mathsf{CT}_\ell$ *resulting in a state* $\widetilde{\mathsf{CT}}$*.*

$\mathsf{Del}(\mathsf{CT})\rightarrow\pi$: *takes as input a ciphertext* CT *and outputs a classical certificate* π*.*

$\mathsf{Vrfy}(\mathsf{vk},\pi)\rightarrow\top$ **or** $\bot$: *takes as input a key* vk *and certificate* π*, and outputs* $\top$ *or* $\bot$*.*

We remark that we frequently overload the functionality of the encryption and decryption procedures by allowing both procedures to take multi-bit messages as input, and to generate or decrypt a sequence of quantum ciphertexts bit-by-bit.

Definition 24 (Compactness and full homomorphism)**.** *A homomorphic encryption scheme with certified deletion* $\mathsf{HE}_{\mathsf{CD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *is fully homomorphic if, for any efficienty (in* $\lambda \in \mathbb{N}$*) computable circuit* $C : \{0,1\}^\ell \to \{0,1\}$ *and any set of inputs* $x = (x_1, \ldots, x_\ell) \in \{0,1\}^\ell$*, it holds that*

$$\Pr\left[\mathsf{Dec}(\mathsf{sk}, \widetilde{\mathsf{CT}}) \neq C(x_1, \ldots, x_\ell) \;\middle|\; \begin{smallmatrix}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KeyGen}(1^\lambda)\\(\mathsf{vk},\mathsf{CT})\leftarrow\mathsf{Enc}(\mathsf{pk},x)\\ \widetilde{\mathsf{CT}}\leftarrow\mathsf{Eval}(C,\mathsf{CT},\mathsf{pk})\end{smallmatrix}\right] \leq \mathrm{negl}(\lambda).$$

We say that a fully homomorphic encryption scheme with certified deletion $(\mathsf{FHE}_{\mathsf{CD}})$ *is compact if its decryption circuit is independent of the circuit* C*. The scheme is leveled fully homomorphic if it takes* 1^L *as an additional input for its key generation procedure and can only evaluate depth* L *Boolean circuits.*

Definition 25 (Correctness of verification)**.** *A homomorphic encryption scheme with certified deletion* $\mathsf{HE}_{\mathsf{CD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *has correctness of verification if the following property holds for any integer* $\lambda \in \mathbb{N}$ *and any set of inputs* $x = (x_1, \ldots, x_\ell) \in \{0,1\}^\ell$

$$\Pr\left[\mathsf{Vrfy}(\mathsf{vk}, \pi) = \bot \;\middle|\; \begin{smallmatrix}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KeyGen}(1^\lambda)\\(\mathsf{vk},\mathsf{CT})\leftarrow\mathsf{Enc}(\mathsf{pk},x)\\ \pi\leftarrow\mathsf{Del}(\mathsf{CT})\end{smallmatrix}\right] \leq \mathrm{negl}(\lambda).$$

Recall that a fully homomorphic encryption scheme with certified deletion enables an untrusted quantum server to compute on encrypted data and to also prove data deletion to a client. In this context, it is desirable for the client to be able to *extract* (i.e., to decrypt) the outcome of the computation without irreversibly affecting the ability of the server to later prove deletion. We use the following definition.

Definition 26 (Extractable FHE scheme with certified deletion)**.** *A fully homomorphic encryption scheme with certified deletion* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Extract}, \mathsf{Del}, \mathsf{Vrfy})$ *is called extractable, if*

- $\mathsf{Eval}(C, \mathsf{CT}_1, \ldots, \mathsf{CT}_\ell, \mathsf{pk})$ *additionally outputs a circuit transcript* t_C *besides* $\widetilde{\mathsf{CT}}$*;*
- $\mathsf{Extract}\langle \mathcal{S}(\varrho, t_C), \mathcal{R}(\mathsf{sk})\rangle$ *is an interactive protocol between a sender* $\mathcal{S}$ *(which takes as input a state* ϱ *and a circuit transcript* t_C*) and a receiver* $\mathcal{R}$ *(which takes as input a key* sk*) with the property that, once the protocol is complete,* $\mathcal{S}$ *obtains a state* $\widetilde{\varrho}$ *and* $\mathcal{R}$ *obtains a bit* $y \in \{0,1\}$*;*

such that for any efficiently computable circuit $C : \{0,1\}^\ell \to \{0,1\}$ *of depth* L *and any input* $x \in \{0,1\}^\ell$*:*

$$\Pr\left[y \neq C(x_1, \ldots, x_\ell) \;\middle|\; \begin{smallmatrix}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KeyGen}(1^\lambda,1^L)\\(\mathsf{vk},\mathsf{CT})\leftarrow\mathsf{Enc}(\mathsf{pk},x)\\(\widetilde{\mathsf{CT}},t_C)\leftarrow\mathsf{Eval}(C,\mathsf{CT},\mathsf{pk})\\(\widetilde{\varrho},y)\leftarrow\mathsf{Extract}\langle\mathcal{S}(\widetilde{\mathsf{CT}},t_C),\mathcal{R}(\mathsf{sk})\rangle\end{smallmatrix}\right] \leq \mathrm{negl}(\lambda), \quad \textit{and}$$

$$\Pr\left[\mathsf{Vrfy}(\mathsf{vk}, \pi) = \bot \;\middle|\; \begin{smallmatrix}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KeyGen}(1^\lambda,1^L)\\(\mathsf{vk},\mathsf{CT})\leftarrow\mathsf{Enc}(\mathsf{pk},x)\\(\widetilde{\mathsf{CT}},t_C)\leftarrow\mathsf{Eval}(C,\mathsf{CT},\mathsf{pk})\\(\widetilde{\varrho},y)\leftarrow\mathsf{Extract}\langle\mathcal{S}(\widetilde{\mathsf{CT}},t_C),\mathcal{R}(\mathsf{sk})\rangle\\ \pi\leftarrow\mathsf{Del}(\widetilde{\varrho})\end{smallmatrix}\right] \leq \mathrm{negl}(\lambda).$$

Remark (Compactness of an extractable FHE scheme)**.** *Our notion of an extractable* FHE *scheme with certified deletion in Definition 26 requires the evaluator to keep a transcript of the circuit that is being applied, which at first sight seems to violate the usual notion of compactness in Definition 24. However, the action of the decryptor during the interactive protocol* Extract *is still independent of the circuit that is being applied, and so it is possible to recover an analogous form of compactness as before.*

8.2 Certified deletion security

Our notion of certified deletion security for homomorphic encryption (HE) schemes is similar to the notion of IND-CPA-CD security for public-key encryption schemes in Definition 22.

Definition 27 (Certified deletion security for HE)**.** *Let* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *be a homomorphic encryption scheme with certified deletion and let* $\mathcal{A}$ *be a* QPT *adversary. We define the security experiment* $\mathsf{Exp}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(b)$ *between* $\mathcal{A}$ *and a challenger as follows:*

1. *The challenger generates a pair* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$*, and sends* pk *to* $\mathcal{A}$*.*
2. $\mathcal{A}$ *sends a distinct plaintext pair* $(m_0, m_1) \in \{0,1\}^\ell \times \{0,1\}^\ell$ *to the challenger.*
3. *The challenger computes* $(\mathsf{vk}, \mathsf{CT}_b) \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b)$*, and sends* $|\mathsf{CT}_b\rangle$ *to* $\mathcal{A}$*.*
4. *At some point in time,* $\mathcal{A}$ *sends a certificate* π *to the challenger.*
5. *The challenger computes* $\mathsf{Vrfy}(\mathsf{vk}, \pi)$ *and sends* sk *to* $\mathcal{A}$*, if the output is* 1*, and* 0 *otherwise.*
6. $\mathcal{A}$ *outputs a guess* $b' \in \{0,1\}$*, which is also the output of the experiment.*

We say that the scheme Σ *is* IND-CPA-CD*-secure if, for any* QPT *adversary* $\mathcal{A}$*, that*

$$\mathsf{Adv}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{Exp}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

9 Dual-Regev Fully Homomorphic Encryption with Certified Deletion

In this section, we describe the main result of this work. We introduce a protocol that allows an untrusted quantum server to perform homomorphic operations on encrypted data, and to simultaneously prove data deletion to a client. Our FHE scheme with certified deletion supports the evaluation of polynomial-sized Boolean circuits composed entirely of NAND gates (see Figure 4) – an assumption we can make without loss of generality, since the NAND operation is universal for classical computation. Note that, for $a, b \in \{0,1\}$, the logical NOT-AND (NAND) operation is defined by

$$\mathsf{NAND}(a,b) = \overline{a \wedge b} = 1 - a \cdot b.$$

Recall also that a Boolean circuit with input $x \in \{0,1\}^n$ is a directed acyclic graph $G = (V, E)$ in which

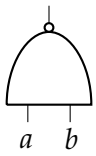

Figure 4: NAND gate.

each node in V is either an input node (corresponding to an input bit x_i), an AND ($\wedge$) gate, an OR ($\vee$) gate, or a NOT ($\neg$) gate. We can naturally identify a Boolean circuit with a function $f : \{0,1\}^n \to \{0,1\}$ which it computes. Due to the universality of the NAND operation, we can represent every Boolean circuit (and the function it computes) with an equivalent circuit consisting entirely of NAND gates. In Figure 5, we give an example of a Boolean circuit composed of three NAND gates that takes as input a string $x \in \{0,1\}^4$.

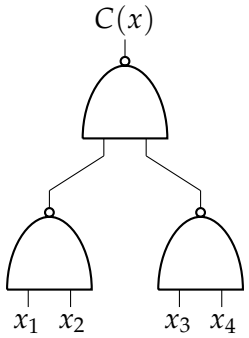

Figure 5: A Boolean circuit C made up of three NAND gates which takes as input a binary string of the form $x \in \{0,1\}^4$. The top-most NAND gate is the designated output node with outcome $C(x) \in \{0,1\}$.

9.1 Construction

In this section, we describe our fully homomorphic encryption scheme with certified deletion. In order to define our construction, we require a so-called *flattening* operation first introduced by Gentry, Sahai and Waters [GSW13] in the context of homomorphic encryption and is also featured in the Dual-Regev FHE scheme of Mahadev [Mah18]. Let $n \in \mathbb{N}$, $q \geq 2$ be a prime modulus and $m \geq 2n \log q$. We define a linear operator $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ called the *gadget matrix*, where $N = (n+1) \cdot \lceil \log q \rceil$. The operator $\mathbf{G}$ converts a binary representation of a vector back to its original vector representation over the ring $\mathbb{Z}_q$. More precisely, for any binary vector $\mathbf{a} = (a_{1,0}, \ldots, a_{1,l-1}, \ldots, a_{m+1,0}, \ldots, a_{m+1,l-1})$ of length N with $\ell = \lceil \log q \rceil$, the matrix $\mathbf{G}$ produces a vector in $\mathbb{Z}_q^{m+1}$ as follows:

$$\mathbf{G}(\mathbf{a}) = \left(\sum_{j=0}^{\lceil \log q \rceil - 1} 2^j \cdot a_{1,j} \ , \ldots, \ \sum_{j=0}^{\lceil \log q \rceil - 1} 2^j \cdot a_{m+1,j} \right). \tag{18}$$

We also define the associated (non-linear) inverse operation $\mathbf{G}^{-1}$ which converts a vector $\mathbf{a} \in \mathbb{Z}_q^{m+1}$ to its binary representation in $\{0,1\}^N$. In other words, we have that $\mathbf{G}^{-1} \cdot \mathbf{G} = \mathbb{1}$ acts as the identity operation.

Our (leveled) FHE scheme with certified deletion is based on the (leveled) Dual-Regev FHE scheme introduced by Mahadev [Mah18] which is a variant of the LWE-based FHE scheme proposed by Gentry, Sahai and Waters [GSW13]. We base our choice of parameters on the aforementioned two works.

Let us first recall the Dual-Regev FHE scheme below.

Construction 2 (Dual-Regev leveled FHE)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. The Dual-Regev leveled* FHE *scheme* $\mathsf{DualFHE} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ *consists of the following* PPT *algorithms:*

$\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: *sample a uniformly random matrix $\bar{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and vector $\bar{\mathbf{x}} \xleftarrow{\$} \{0,1\}^m$ and let $\mathbf{A} = [\bar{\mathbf{A}} | \bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q]^T$. Output $(\mathsf{pk}, \mathsf{sk})$, where $\mathsf{pk} = \mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}$ and $\mathsf{sk} = (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$.*

$\mathsf{Enc}(\mathsf{pk}, x)$: *to encrypt $x \in \{0,1\}$, parse $\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n} \leftarrow \mathsf{pk}$, sample $\mathbf{S} \xleftarrow{\$} \mathbb{Z}_q^{n\times N}$ and $\mathbf{E} \sim D_{\mathbb{Z}^{(m+1)\times N}, \alpha q}$ and output $\mathsf{CT} = \mathbf{A} \cdot \mathbf{S} + \mathbf{E} + x \cdot \mathbf{G} \pmod q \in \mathbb{Z}_q^{(m+1)\times N}$, where $\mathbf{G}$ is the gadget matrix in Eq. (18).*

$\mathsf{Eval}(C, \mathsf{CT})$: *apply the circuit C composed of* NAND *gates on a ciphertext tuple* CT *as follows:*

- *parse the ciphertext tuple as* $(\mathsf{CT}_1, \ldots, \mathsf{CT}_\ell) \leftarrow \mathsf{CT}$.
- *repeat for every* NAND *gate in* C*: to apply a* NAND *gate on a ciphertext pair* $(\mathsf{CT}_i, \mathsf{CT}_j)$*, parse matrices* $\mathbf{C}_i \leftarrow \mathsf{CT}_i$ *and* $\mathbf{C}_j \leftarrow \mathsf{CT}_j$ *with* $\mathbf{C}_i, \mathbf{C}_j \in \mathbb{Z}_q^{(m+1)\times N}$ *and generate*

$$\mathbf{C}_{ij} = \mathbf{G} - \mathbf{C}_i \cdot \mathbf{G}^{-1}(\mathbf{C}_j) \pmod q.$$

  Let $\mathsf{CT}_{ij} \leftarrow \mathbf{C}_{ij}$ *denote the outcome ciphertext.*

$\mathsf{Dec}(\mathsf{sk}, \mathsf{CT})$: *parse* $\mathbf{C} \in \mathbb{Z}_q^{(m+1)\times N} \leftarrow \mathsf{CT}$ *and compute* $c = \mathsf{sk}^T \cdot \mathbf{c}_N \in \mathbb{Z} \cap (-\frac{q}{2}, \frac{q}{2}]$*, where* $\mathbf{c}_N \in \mathbb{Z}_q^{m+1}$ *is the* N*-th column of* $\mathbf{C}$*, and then output* 0*, if* c *is closer to* 0 *than to* $\lfloor \frac{q}{2} \rfloor$*, and output* 1*, otherwise.*

The Dual-Regev FHE scheme supports the homomorphic evaluation of a NAND gate in the following sense. If CT_0 and CT_1 are ciphertexts that encrypt two bits x_0 and x_1, respectively, then the resulting outcome $\mathsf{CT} = \mathbf{G} - \mathsf{CT}_0 \cdot \mathbf{G}^{-1}(\mathsf{CT}_1) \pmod q$ is an encryption of $\mathsf{NAND}(x_0, x_1) = 1 - x_0 \cdot x_1$, where $\mathbf{G}$ is the gadget matrix that converts a binary representation of a vector back to its original representation over the ring $\mathbb{Z}_q$. Moreover, the new ciphertext CT maintains the form of an LWE sample with respect to the same public key pk, albeit for a new LWE secret and a new (non-necessarily Gaussian) noise term of bounded magnitude. This property is crucial, as knowledge of the secret key sk (i.e., a short trapdoor vector) still allows for the decryption of the ciphertext CT once a NAND gate has been applied.

The following result is implicit in the work of Mahadev [Mah18, Theorem 5.1].

Theorem 8 ([Mah18])**.** *Let* $\lambda \in \mathbb{N}$ *be the security parameter. Let* $n \in \mathbb{N}$*, let* $q \geq 2$ *be a prime modulus and* $m \geq 2n \log q$*. Let* $N = (n+1) \cdot \lceil \log q \rceil$ *be an integer and let* L *be an upper bound on the depth of the polynomial-sized Boolean circuit which is to be evaluated. Let* $\alpha \in (0,1)$ *be a ratio such that*

$$2\sqrt{n} \leq \alpha q \leq \frac{q}{4(m+1) \cdot N \cdot (N+1)^L}.$$

Then, the scheme in Construction 2 is an IND-CPA*-secure leveled fully homomorphic encryption scheme under the* $\mathsf{LWE}_{n,q,\alpha q}^{(m+1)\times N}$ *assumption.*

Note that the Dual-Regev FHE scheme is *leveled* in the sense that an apriori upper bound L on the NAND-depth of the circuit is required to set the parameters appropriately. We remark that a proper (non-leveled) FHE scheme can be obtained under an additional circular security assumption [BV11].

The leveled Dual-Regev FHE scheme inherits a crucial property from its public-key counterpart. Namely, in contrast to the FHE scheme in [GSW13], the ciphertext takes the form of a regular sample from the LWE distribution together with an additive shift $x \cdot \mathbf{G}$ that depends on the plaintext $x \in \{0,1\}$. In particular, if a Boolean circuit C of polynomial NAND-depth L is applied to the ciphertext corresponding to a plaintext $x \in \{0,1\}^\ell$ in Construction 2, then the resulting final ciphertext is of the form $\mathbf{A} \cdot \mathbf{S} + \mathbf{E} + C(x)\mathbf{G}$, where $\mathbf{S} \in \mathbb{Z}_q^{n\times N}$, $\mathbf{E} \in \mathbb{Z}_q^{(m+1)\times N}$ and $\|\mathbf{E}\|_\infty \leq \alpha q \sqrt{(m+1)N} \cdot (N+1)^L$ (see [GSW13] for details). Choosing $1/\alpha$ to be sub-exponential in N as in [GSW13], we can therefore allow for homomorphic computations of arbitrary polynomial-sized Boolean circuits of NAND-depth at most L. It is easy to see that the decryption procedure of the leveled Dual-Regev FHE scheme is successful as long as the cumulative error $\mathbf{E}$ satisfies the condition $\|\mathbf{E}\|_\infty \leq \frac{q}{4\sqrt{(m+1)N}}$.

This property is essential as it allows us to extend Dual-Regev PKE scheme with certified deletion towards a leveled FHE scheme, which we denote by $\mathsf{FHE}_{\mathsf{CD}}$. Using Gaussian coset states, we can again encode Dual-Regev ciphertexts for the purpose of certified deletion while simultaneously preserving their cryptographic functionality.

Dual-Regev leveled FHE with certified deletion. Let us now describe our (leveled) FHE scheme with certified deletion. We base our choice of parameters on the Dual-Regev FHE scheme of Mahadev [Mah18] which is a variant of the scheme due to Gentry, Sahai and Waters [GSW13].

Parameters. Let $\lambda \in \mathbb{N}$ be the security parameter and let $n \in \mathbb{N}$. Let L be an upper bound on the depth of the polynomial-sized Boolean circuit which is to be evaluated. We choose the following set of parameters for the Dual-Regev leveled FHE scheme (each parameterized by the security parameter λ).

- a prime modulus $q \geq 2$.
- an integer $m \geq 2n \log q$.
- an integer $N = (n+1) \cdot \lceil \log q \rceil$.
- a noise ratio $\alpha \in (0,1)$ such that

$$\sqrt{8(m+1)N} \leq \alpha q \leq \frac{q}{\sqrt{8}(m+1) \cdot N \cdot (N+1)^L}.$$

Construction 3 (Dual-Regev leveled FHE scheme with certified deletion)**.** *Let $\lambda \in \mathbb{N}$ be a parameter and* $\mathsf{DualFHE} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ *be the scheme in Construction 2. The Dual-Regev (leveled)* FHE *scheme* $\mathsf{DualFHE}_{\mathsf{CD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *with certified deletion is defined by:*

$\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: *generate* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{DualFHE.KeyGen}(1^\lambda)$ *and output* $(\mathsf{pk}, \mathsf{sk})$.

$\mathsf{Enc}(\mathsf{pk}, x) \to (\mathsf{vk}, |\mathsf{CT}\rangle)$: *to encrypt a bit $x \in \{0,1\}$, parse $\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n} \leftarrow \mathsf{pk}$ and, for $i \in [N]$, run $(|\psi_{\mathbf{y}_i}\rangle, \mathbf{y}_i) \leftarrow \mathsf{GenPrimal}(\mathbf{A}^T, 1/\alpha)$ in Algorithm 2, where $\mathbf{y}_i \in \mathbb{Z}_q^n$, and output the pair*

$$\left(\mathsf{vk} \leftarrow (\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}, (\mathbf{y}_1, \ldots, \mathbf{y}_N) \in \mathbb{Z}_q^{n\times N}), \quad |\mathsf{CT}\rangle \leftarrow \mathbf{X}_q^{x\cdot \mathbf{g}_1} |\psi_{\mathbf{y}_1}\rangle \otimes \cdots \otimes \mathbf{X}_q^{x\cdot \mathbf{g}_N} |\psi_{\mathbf{y}_N}\rangle\right),$$

where $(\mathbf{g}_1, \ldots, \mathbf{g}_N)$ are the rows of the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ in Eq. (18).

$\mathsf{Eval}(C, |\mathsf{CT}\rangle) \to (|\widetilde{\mathsf{CT}}\rangle, t_C)$: *apply the Boolean circuit C composed of* NAND *gates to the ciphertext $|\mathsf{CT}\rangle$ in system $\mathsf{C}_{\mathsf{in}} = \mathsf{C}_1 \cdots \mathsf{C}_\ell$ as follows: For every gate NAND_{ij} in the circuit C between a ciphertext pair in systems C_i and C_j, repeat the following two steps:*

- *apply U_{NAND} from Definition 28 to systems $\mathsf{C}_i\mathsf{C}_j$ of the ciphertext CT by appending an auxiliary system C_{ij}. This results in a new ciphertext state CT which contains the additional system C_{ij}.*
- *add the gate NAND_{ij} to the circuit transcript t_C.*

Output $(|\widetilde{\mathsf{CT}}\rangle, t_C)$, where $|\widetilde{\mathsf{CT}}\rangle$ is the final post-evaluation state in systems $\mathsf{C}_{\mathsf{in}}\mathsf{C}_{\mathsf{aux}}\mathsf{C}_{\mathsf{out}}$ and

- *$\mathsf{C}_{\mathsf{in}} = \mathsf{C}_1 \cdots \mathsf{C}_\ell$ denotes the initial ciphertext systems of $|\mathsf{CT}_1\rangle \otimes \cdots \otimes |\mathsf{CT}_\ell\rangle$.*
- *$\mathsf{C}_{\mathsf{aux}}$ denotes all intermediate auxiliary ciphertext systems.*
- *$\mathsf{C}_{\mathsf{out}}$ denotes the final ciphertext system corresponding to the output of the circuit C.*

$\mathsf{Dec}(\mathsf{sk}, |\mathsf{CT}\rangle) \to \{0,1\}^\mu$ **or** $\perp$: *measure the ciphertext $|\mathsf{CT}\rangle$ in the computational basis to obtain an outcome $\mathbf{C}$ and output $x' \leftarrow \mathsf{DualFHE.Dec}(\mathsf{sk}, \mathbf{C})$.*

$\mathsf{Del}(|\mathsf{CT}\rangle) \to \pi$: *measure* $|\mathsf{CT}\rangle$ *in the Fourier basis with outcomes* $\pi = (\pi_1, \ldots, \pi_N) \in \mathbb{Z}_q^{(m+1)\times N}$.

$\mathsf{Extract}\langle \mathcal{S}(|\widetilde{\mathsf{CT}}\rangle, t_C), \mathcal{R}(\mathsf{sk})\rangle \to (\varrho, y)$ *this is the following interactive protocol between a sender* $\mathcal{S}$ *with input* $|\widetilde{\mathsf{CT}}\rangle$ *in systems* $\mathsf{C_{in}C_{aux}C_{out}}$ *and transcript* t_C*, and a receiver* $\mathcal{R}$ *with input* sk*:*

- $\mathcal{S}$ *and* $\mathcal{R}$ *run the rewinding protocol* $\Pi = \langle \mathcal{S}(|\widetilde{\mathsf{CT}}\rangle, t_C), \mathcal{R}(\mathsf{sk})\rangle$ *in Protocol 1.*
- *Once* Π *is complete,* $\mathcal{S}$ *obtains a state* ϱ *in system* $\mathsf{C_{in}}$ *and* $\mathcal{R}$ *obtains a bit* $y \in \{0,1\}$.

$\mathsf{Vrfy}(\mathsf{vk}, \mathsf{pk}, \pi) \to \{0,1\}$: *to verify the deletion certificate* $\pi = (\pi_1, \ldots, \pi_N) \in \mathbb{Z}_q^{(m+1)\times N}$*, parse* $(\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}, (\mathbf{y}_1, \ldots, \mathbf{y}_N) \in \mathbb{Z}_q^{n\times N}) \leftarrow \mathsf{vk}$ *and output* $\top$*, if* $\mathbf{A}^T \cdot \pi_i = \mathbf{y}_i \pmod q$ *and* $\|\pi_i\| \leq \sqrt{m+1}/\sqrt{2}\alpha$ *for every* $i \in [N]$*, and output* $\bot$*, otherwise.*

Protocol 1 (Rewinding Protocol)**.** *Let* $\mathsf{DualFHE} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ *be the Dual-Regev* FHE *scheme in Construction 2. Consider the following interactive protocol* $\Pi = \langle \mathcal{S}(\varrho, t_C), \mathcal{R}(\mathsf{sk})\rangle$ *between a sender* $\mathcal{S}$ *which takes as input state* ϱ *in systems* $\mathsf{C_{in}C_{aux}C_{out}}$ *and a transcript* t_C *of a Boolean circuit* C*, as well as a receiver* $\mathcal{R}$ *which takes as input a secret key* sk*.*

1. $\mathcal{S}$ *sends system* $\mathsf{C_{out}}$ *of the state* ϱ *associated with the encrypted output of the circuit* C *to* $\mathcal{R}$*.*
2. $\mathcal{R}$ *runs* $U_{\mathsf{DualFHE.Dec_{sk}}}$ *(with the key* sk *hard coded) to reversibly decrypt system* $\mathsf{C_{out}}$*, where*

$$U_{\mathsf{DualFHE.Dec_{sk}}} : \quad |\mathbf{C}\rangle_{\mathsf{C_{out}}} \otimes |0\rangle_M \to |\mathbf{C}\rangle_{\mathsf{C_{out}}} \otimes |\mathsf{DualFHE.Dec_{sk}}(\mathbf{C})\rangle_M ,$$

   for any matrix $\mathbf{C} \in \mathbb{Z}_q^{(m+1)\times N}$*.* $\mathcal{R}$ *then measures system* M *to obtain a bit* $y \in \{0,1\}$ *(the supposed output of the Boolean circuit* C*). Afterwards,* $\mathcal{R}$ *applies* $U^\dagger_{\mathsf{DualFHE.Dec_{sk}}}$*, discards the ancillary system* M*, and sends back the post-measurement system* $\widetilde{\mathsf{C_{out}}}$ *of the resulting ciphertext* $\widetilde{\varrho}$ *to* $\mathcal{S}$*.*
3. $\mathcal{S}$ *repeats the following two steps in order to uncompute the systems* $\mathsf{C_{aux}}\widetilde{\mathsf{C_{out}}}$ *from the state* $\widetilde{\varrho}$*: For every gate* $\mathsf{NAND}_{ij} \in t_C$*, where* i *and* j *denote the respective ciphertext systems* C_i *and* C_j*, in decreasing order starting from the last gate in the circuit transcript* t_C*:*
   - $\mathcal{S}$ *applies* $U^\dagger_{\mathsf{NAND}}$ *from Definition 28 to systems* $\mathsf{C}_i\mathsf{C}_j\mathsf{C}_{ij}$ *of* $\widetilde{\varrho}$ *to uncompute system* C_{ij}*.*
   - $\mathcal{S}$ *repeats the procedure starting from the new outcome state* $\widetilde{\varrho}$*.*

Let us now define how to perform the homomorphic NAND gate in Construction 3 in more detail.

Definition 28 (Homomorphic NAND gate)**.** *Let* $q \geq 2$ *be a modulus, and let* m *and* N *be integers. Let* $\mathbf{X}, \mathbf{Y}, \mathbf{Z} \in \mathbb{Z}_q^{(m+1)\times N}$ *be arbitrary matrices. We define the homomorphic* NAND *gate as the unitary*

$$U_{\mathsf{NAND}} : \quad |\mathbf{X}\rangle_X \otimes |\mathbf{Y}\rangle_Y \otimes |\mathbf{Z}\rangle_Z \quad \to \quad |\mathbf{X}\rangle_X \otimes |\mathbf{Y}\rangle_Y \otimes |\mathbf{Z} + \mathbf{G} - \mathbf{X}\cdot\mathbf{G}^{-1}(\mathbf{Y}) \pmod q\rangle_Z ,$$

where $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ *is the gadget matrix in Eq. (18).*

To illustrate the action of our homomorphic NAND gate, we consider a simple example.

Example. Consider a pair of two ciphertexts $|\mathsf{CT}_i\rangle \otimes |\mathsf{CT}_j\rangle$ which encrypt two bits $x_i, x_j \in \{0,1\}$ as in Construction 3. Let $U_{\mathsf{NAND}_{ij}}$ denote the homomorphic NAND gate applied to systems C_i and C_j. Then,

$$U_{\mathsf{NAND}_{ij}} : \quad |\mathsf{CT}_i\rangle_{C_i} \otimes |\mathsf{CT}_j\rangle_{C_j} \otimes |\mathbf{0}\rangle_{C_{ij}} \quad \to \quad |\mathsf{CT}_{ij}\rangle_{C_iC_jC_{ij}}.$$

Here, $|\mathsf{CT}_{ij}\rangle$ is the resulting ciphertext in systems $C_iC_jC_{ij}$. Note that $U_{\mathsf{NAND}_{ij}}$ is reversible in the sense that

$$U^\dagger_{\mathsf{NAND}_{ij}} : \quad |\mathsf{CT}_{ij}\rangle_{C_iC_jC_{ij}} \quad \to \quad |\mathsf{CT}_i\rangle_{C_i} \otimes |\mathsf{CT}_j\rangle_{C_j} \otimes |\mathbf{0}\rangle_{C_{ij}}.$$

Let us now analyze how U_{NAND} acts on the basis states of a pair of ciphertexts $|\mathsf{CT}_i\rangle \otimes |\mathsf{CT}_j\rangle$ that encode LWE samples as in Construction 3. In the following, $\mathbf{E}_i, \mathbf{E}_j \sim D_{\mathbb{Z}_q^{(m+1)\times N}, \frac{\alpha q}{\sqrt{2}}}$ have a (truncated) discrete Gaussian distribution as part of the superposition. Then,

$$\begin{aligned} U_{\mathsf{NAND}_{ij}} : \quad & |\mathbf{A}\mathbf{S}_i + \mathbf{E}_i + x_i\mathbf{G}\rangle_{C_i} \otimes |\mathbf{A}\mathbf{S}_j + \mathbf{E}_j + x_j\mathbf{G}\rangle_{C_j} \otimes |\mathbf{0}\rangle_{C_{ij}} \\ \to \; & |\mathbf{A}\mathbf{S}_i + \mathbf{E}_i + x_i\mathbf{G}\rangle_{C_i} \otimes |\mathbf{A}\mathbf{S}_j + \mathbf{E}_j + x_j\mathbf{G}\rangle_{C_j} \otimes |\mathbf{A}\mathbf{S}_{ij} + \mathbf{E}_{ij} + (1 - x_ix_j)\mathbf{G}\rangle_{C_{ij}}, \end{aligned}$$

where introduced the following matrices

$$\begin{aligned} \mathbf{S}_{ij} &:= -\mathbf{S}_i \cdot \mathbf{G}^{-1}(\mathbf{A}\mathbf{S}_j + \mathbf{E}_j + x_j\mathbf{G}) - x_i\mathbf{S}_i \pmod q \\ \mathbf{E}_{ij} &:= -\mathbf{E}_i \cdot \mathbf{G}^{-1}(\mathbf{A}\mathbf{S}_j + \mathbf{E}_j + x_j\mathbf{G}) - x_i\mathbf{E}_j \pmod q. \end{aligned}$$

Because the initial error terms have the property that $\|\mathbf{E}_i\|_\infty, \|\mathbf{E}_j\|_\infty \leq \alpha q\sqrt{(m+1)N/2}$, it follows that the resulting error after a single NAND gate is at most (see also [GSW13, Mah18] for more details)

$$\|\mathbf{E}_{ij}\|_\infty \leq \alpha q\sqrt{\frac{(m+1)N}{2}} \cdot (N+1).$$

In other words, the cumulative error term remains short relative to the modulus q after every application of a homomorphic NAND gate, exactly as in the Dual-Regev FHE scheme of Mahadev [Mah18].

9.2 Rewinding lemma

Notice that the procedure $\mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Eval}$ in Construction 3 produces a highly entangled state since the unitary operation U_{NAND} induces entanglement between the Gaussian noise terms. In the next lemma, we show that it is possible to *rewind* the evaluation procedure to be able to prove data deletion to a client.

Lemma 19 (Rewinding lemma)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n \in \mathbb{N}$, let $q \geq 2$ be a prime modulus and $m \geq 2n\log q$. Let $N = (n+1)\cdot\lceil\log q\rceil$ be an integer and let L be an upper bound on the depth of the polynomial-sized Boolean circuit which is to be evaluated. Let $\alpha \in (0,1)$ be a ratio such that*

$$\sqrt{8(m+1)N} \leq \alpha q \leq \frac{q}{\sqrt{8}(m+1)\cdot N\cdot(N+1)^L}.$$

Let $\mathsf{DualFHE}_{\mathsf{CD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ be the Dual-Regev (leveled) FHE scheme with certified deletion in Construction 3 and let Π be the interactive protocol in Protocol 1. Then, the following holds for any parameter $\lambda \in \mathbb{N}$, plaintext $x \in \{0,1\}^\ell$ and any polynomial-sized Boolean circuit C:

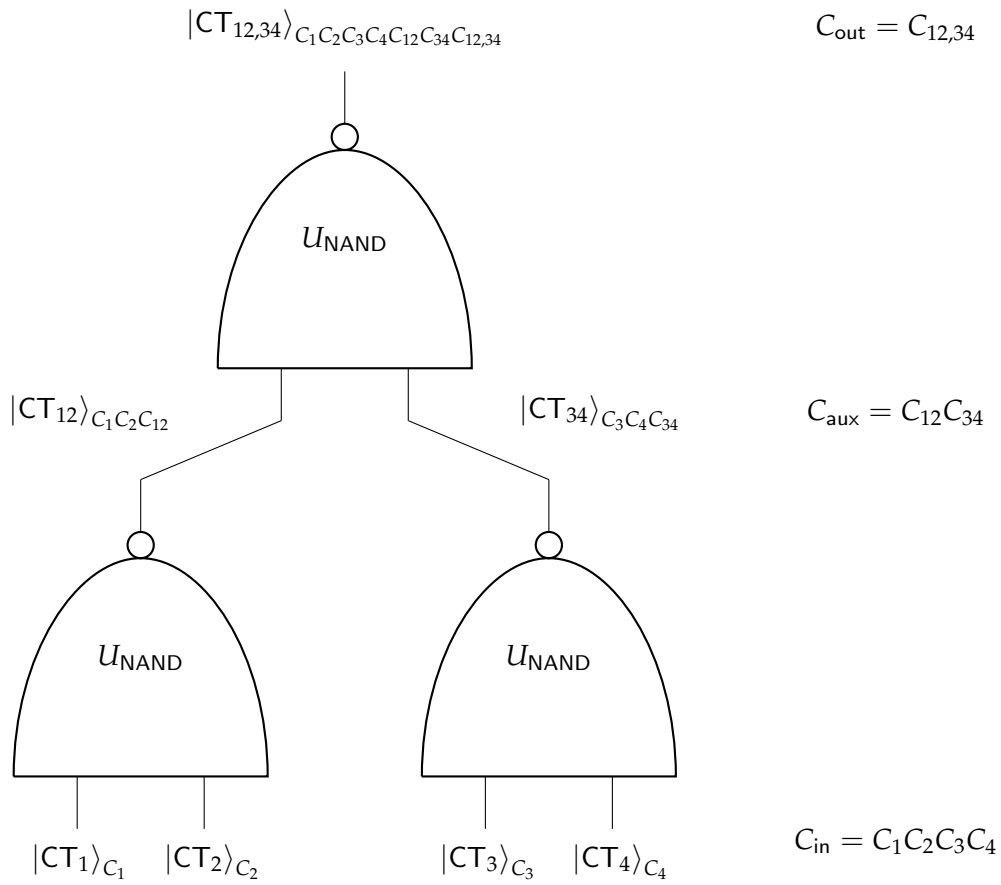

Figure 6: Homomorphic evaluation of a Boolean circuit C composed entirely of three NAND gates. Here, the input is the quantum ciphertext $|\mathsf{CT}_1\rangle \otimes |\mathsf{CT}_2\rangle \otimes |\mathsf{CT}_3\rangle \otimes |\mathsf{CT}_4\rangle$ which corresponds to an encryption of the plaintext $x = (x_1, \ldots, x_4) \in \{0,1\}^4$ as in Construction 3. The resulting ciphertext $|\mathsf{CT}_{12,34}\rangle$ lives on a system $C_1C_2C_3C_4C_{12}C_{34}C_{12,34}$ of which the last system $C_{12,34}$ contains an encryption of $C(x) \in \{0,1\}$.

After the interactive protocol $\Pi = \langle \mathcal{S}(|\widetilde{\mathsf{CT}}\rangle, t_C), \mathcal{R}(\mathsf{sk})\rangle$ *between the sender* $\mathcal{S}$ *and receiver* $\mathcal{R}$ *is complete, the sender* $\mathcal{S}$ *is in possession of a quantum state* ϱ *in system* C_{in} *that satisfies*

$$\|\varrho - |\mathsf{CT}\rangle\langle\mathsf{CT}|\|_{\mathrm{tr}} \leq \mathrm{negl}(\lambda),$$

where $(|\widetilde{\mathsf{CT}}\rangle, t_C) \leftarrow \mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Eval}(C, |\mathsf{CT}\rangle)$ *is the post-evaluation state* $|\widetilde{\mathsf{CT}}\rangle$ *in systems* $C_{\mathsf{in}}C_{\mathsf{aux}}C_{\mathsf{out}}$ *and where* $|\mathsf{CT}\rangle \leftarrow \mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Enc}(\mathsf{pk}, x)$ *is the initial state for* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{KeyGen}(1^\lambda)$.

Proof. Let $\lambda \in \mathbb{N}$, $x \in \{0,1\}^\ell$ be a plaintext and C be any Boolean circuit of NAND-depth $L = \mathrm{poly}(\lambda)$. Let $(|\widetilde{\mathsf{CT}}\rangle, t_C) \leftarrow \mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Eval}(C, |\mathsf{CT}\rangle)$ be the post-evaluation state $|\widetilde{\mathsf{CT}}\rangle$ in systems $C_{\mathsf{in}}C_{\mathsf{aux}}C_{\mathsf{out}}$ with circuit transcript t_C and let ϱ be the outcome of the interactive protocol $\Pi = \langle \mathcal{S}(|\widetilde{\mathsf{CT}}\rangle, t_C), \mathcal{R}(\mathsf{sk})\rangle$. Recall that, in Lemma 20, we established that there exists a negligible $\varepsilon(\lambda)$ such that $\mathsf{DualFHE}.\mathsf{Dec}_{\mathsf{sk}}$ decrypts system C_{out} of $|\widetilde{\mathsf{CT}}\rangle$ with probability at least $1 - \varepsilon$. By the "Almost As Good As New Lemma" (Lemma 1), performing the operation $U_{\mathsf{DualFHE}.\mathsf{Dec}_{\mathsf{sk}}}$, measuring the ancillary register M and rewinding the computation, results in a mixed state $\widetilde{\varrho}$ that is within trace distance $\sqrt{\varepsilon}$ of the post-evaluation state $|\widetilde{\mathsf{CT}}\rangle$. Notice that, by reversing the sequence U_{t_C} of homomorphic NAND gates according to the transcript t_C

with respect to $|\widetilde{\mathsf{CT}}\rangle$, we recover the initial ciphertext $|\mathsf{CT}\rangle\langle\mathsf{CT}| = U_{t_C}^\dagger|\widetilde{\mathsf{CT}}\rangle\langle\widetilde{\mathsf{CT}}|\,U_{t_C}$ in system C_{in}. By definition, we also have that $\varrho = U_{t_C}^\dagger\widetilde{\varrho}\,U_{t_C}$. Therefore,

$$\|\varrho - |\mathsf{CT}\rangle\langle\mathsf{CT}|\|_{\mathrm{tr}} = \|U_{t_C}^\dagger\widetilde{\varrho}\,U_{t_C} - U_{t_C}^\dagger|\widetilde{\mathsf{CT}}\rangle\langle\widetilde{\mathsf{CT}}|\,U_{t_C}\|_{\mathrm{tr}} = \|\widetilde{\varrho} - |\widetilde{\mathsf{CT}}\rangle\langle\widetilde{\mathsf{CT}}|\|_{\mathrm{tr}} \leq \sqrt{\varepsilon(\lambda)},$$

where we used that the trace distance is unitarily invariant. Since $\varepsilon(\lambda) = \mathrm{negl}(\lambda)$, this proves the claim. □

Proof of correctness. Let us now verify the correctness of decryption and verification of Construction 3.

Lemma 20 (Compactness and full homomorphism of $\mathsf{DualFHE}_{\mathsf{CD}}$)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n \in \mathbb{N}$, let $q \geq 2$ be a prime and $m \geq 2n\log q$. Let $N = (n+1)\cdot\lceil\log q\rceil$ and let L be an upper bound on the depth of the polynomial-sized Boolean circuit which is to be evaluated. Let $\alpha \in (0,1)$ be a ratio with*

$$\sqrt{8(m+1)N} \leq \alpha q \leq \frac{q}{\sqrt{8}(m+1)\cdot N\cdot(N+1)^L}.$$

Then, the scheme $\mathsf{DualFHE}_{\mathsf{CD}} = (\mathsf{KeyGen},\mathsf{Enc},\mathsf{Dec},\mathsf{Eval},\mathsf{Del},\mathsf{Vrfy})$ in Construction 3 is a compact and fully homomorphic encryption scheme with certified deletion. In other words, for any efficienty (in $\lambda\in\mathbb{N}$) computable circuit $C : \{0,1\}^\ell \to \{0,1\}$ and any set of inputs $x = (x_1,\ldots,x_\ell) \in \{0,1\}^\ell$, it holds that:

$$\Pr\left[\mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Dec}(\mathsf{sk},|\widetilde{\mathsf{CT}}\rangle) \neq C(x_1,\ldots,x_\ell) \;\middle|\; \begin{smallmatrix}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{KeyGen}(1^\lambda,1^L)\\ (\mathsf{vk},|\mathsf{CT}\rangle)\leftarrow\mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Enc}(\mathsf{pk},x)\\ (|\widetilde{\mathsf{CT}}\rangle,t_C)\leftarrow\mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Eval}(C,|\mathsf{CT}\rangle,\mathsf{pk})\end{smallmatrix}\right] \leq \mathrm{negl}(\lambda).$$

Proof. Let $|\mathsf{CT}\rangle$ be the ciphertext output by $\mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Enc}(\mathsf{pk},x)$, where $x\in\{0,1\}^\ell$ denotes the plaintext, and let $(|\widetilde{\mathsf{CT}}\rangle, t_C) \leftarrow \mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Eval}(C,|\mathsf{CT}\rangle)$ be the output of the evaluation procedure. Let us first consider the case when $t_C = \emptyset$, i.e. not a single NAND gate has been applied to the ciphertext. In this case, the claim follows from the fact that the truncated discrete Gaussian $D_{\mathbb{Z}_q^{(m+1)\times N},\frac{\alpha q}{\sqrt{2}}}$ is supported on $\{\mathbf{X}\in\mathbb{Z}_q^{(m+1)\times N} : \|\mathbf{X}\|_\infty \leq \alpha q\sqrt{N(m+1)/2}\}$. Recall that $\mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Dec}(\mathsf{sk},|\widetilde{\mathsf{CT}}\rangle)$ measures the ciphertext $|\widetilde{\mathsf{CT}}\rangle$ in the computational basis with outcome $\mathbf{C} = (\mathbf{C}_1,\ldots,\mathbf{C}_\ell)$, where $\mathbf{C}_i\in\mathbb{Z}_q^{(m+1)\times N}$ is a matrix, and outputs $x'\leftarrow\mathsf{DualFHE}.\mathsf{Dec}(\mathsf{sk},\mathbf{C})$. By our choice of parameters, each error term satisfies

$$\|\mathbf{E}_i\|_\infty \leq \alpha q\sqrt{\frac{N(m+1)}{2}} < \frac{q}{4\sqrt{(m+1)N}}, \quad \forall i\in[\ell].$$

Hence, decryption correctness is preserved if $t_C = \emptyset$. Let us now consider the case when $t_C \neq \emptyset$, i.e. the Boolean circuit C consists of at least one NAND gate which has been applied to the ciphertext $|\mathsf{CT}\rangle$. In this case, the cumulative error in system C_{out} after L applications of U_{NAND} in Definition 28 is at most $\alpha q\sqrt{(m+1)N/2}(N+1)^L$, which is less than $\frac{q}{4\sqrt{(m+1)N}}$ by our choice of parameters. Therefore, the procedure $\mathsf{DualFHE}.\mathsf{Dec}_{\mathsf{sk}}$ decrypts a computational basis state in system C_{out} of the state $|\widetilde{\mathsf{CT}}\rangle$ correctly with probability at least $1-\mathrm{negl}(\lambda)$. Furthermore, because the procedure $\mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Dec}$ is independent of the circuit C and its depth L, the scheme $\mathsf{DualFHE}_{\mathsf{CD}}$ is compact. This proves the claim. □

Let us now verify the correctness of verification of the scheme $\mathsf{DualFHE}_{\mathsf{CD}}$ in Construction 3 according to Definition 25. We show the following.

Lemma 21 (Correctness of verification)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n \in \mathbb{N}$, let $q \geq 2$ be a prime modulus and $m \geq 2n \log q$. Let $N = (n+1) \cdot \lceil \log q \rceil$ be an integer and let L be an upper bound on the depth of the polynomial-sized Boolean circuit which is to be evaluated. Let $\alpha \in (0,1)$ be a ratio with*

$$\sqrt{8(m+1)N} \leq \alpha q \leq \frac{q}{\sqrt{8}(m+1) \cdot N \cdot (N+1)^L}.$$

Then, the Dual-Regev FHE *scheme* $\mathsf{DualFHE}_{\mathsf{CD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *with certified deletion in Construction 3 satisfies verification correctness. In other words, for any $\lambda \in \mathbb{N}$, any plaintext $x \in \{0,1\}^\ell$ and any polynomial-sized Boolean circuit C entirely composed of* NAND *gates:*

$$\Pr\left[\mathsf{Verify}(\mathsf{vk}, \pi) = 1 \;\middle|\; \begin{smallmatrix} (\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ (\mathsf{vk}, |\mathsf{CT}\rangle) \leftarrow \mathsf{Enc}(\mathsf{pk}, x) \\ \pi \leftarrow \mathsf{Del}(|\mathsf{CT}\rangle) \end{smallmatrix}\right] \geq 1 - \mathrm{negl}(\lambda).$$

Proof. Consider a bit $x \in \{0,1\}$ and a public key pk given by $\mathbf{A} = [\bar{\mathbf{A}} | \bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q] \in \mathbb{Z}_q^{(m+1)\times n}$, for $\bar{\mathbf{x}} \xleftarrow{\$} \{0,1\}^m$. By the Leftover Hash Lemma (Lemma 4), the distribution of $\mathbf{A}$ is within negligible total variation distance of the uniform distribution over $\mathbb{Z}_q^{(m+1)\times n}$. Lemma 9 implies that the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$ with overwhelming probability. We consider the ciphertext $|\mathsf{CT}\rangle$ output by $\mathsf{Enc}(\mathsf{pk}, x)$, where

$$|\mathsf{CT}\rangle \leftarrow \mathbf{X}_q^{x \cdot \mathbf{g}_1} |\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes \mathbf{X}_q^{x \cdot \mathbf{g}_N} |\hat{\psi}_{\mathbf{y}_N}\rangle,$$

and where $(\mathbf{g}_1, \ldots, \mathbf{g}_N)$ are the rows of the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ in Eq. (18). Given our choice,

$$\sqrt{8(m+1)N} \leq \alpha q \leq \frac{q}{\sqrt{8}(m+1) \cdot N \cdot (N+1)^L},$$

Corollary 1 implies that the Fourier transform of $|\mathsf{CT}\rangle$ is within negligible trace distance of the state

$$|\widehat{\mathsf{CT}}\rangle = \sum_{\substack{\mathbf{x}_1 \in \mathbb{Z}_q^{m+1}: \\ \mathbf{A}\mathbf{x}_1 = \mathbf{y}_1 \pmod q}} \varrho_{\frac{1}{\alpha}}(\mathbf{x}_1)\, \omega_q^{\langle \mathbf{x}_1, x \cdot \mathbf{g}_1 \rangle} |\mathbf{x}_1\rangle \otimes \cdots \otimes \sum_{\substack{\mathbf{x}_N \in \mathbb{Z}_q^{m+1}: \\ \mathbf{A}\mathbf{x}_N = \mathbf{y}_N \pmod q}} \varrho_{\frac{1}{\alpha}}(\mathbf{x}_N)\, \omega_q^{\langle \mathbf{x}_N, x \cdot \mathbf{g}_N \rangle} |\mathbf{x}_N\rangle.$$

From Lemma 11, it follows that the distribution of computational basis measurement outcomes is within negligible total variation distance of the sample

$$\pi = (\pi_1, \ldots, \pi_N) \sim D_{\Lambda_q^{\mathbf{y}_1}(\mathbf{A}), \frac{1}{\sqrt{2}\alpha}} \times \cdots \times D_{\Lambda_q^{\mathbf{y}_N}(\mathbf{A}), \frac{1}{\sqrt{2}\alpha}},$$

where $\|\pi_i\| \leq \sqrt{m+1}/\sqrt{2}\alpha$ for every $i \in [N]$. This proves the claim. □

We now show that our scheme $\mathsf{DualFHE}_{\mathsf{CD}}$ in Construction 3 is *extractable* according to Definition 26.

Lemma 22 (Extractability of $\mathsf{DualFHE}_{\mathsf{CD}}$)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n \in \mathbb{N}$, let $q \geq 2$ be a prime modulus and $m \geq 2n \log q$. Let $N = (n+1) \cdot \lceil \log q \rceil$ and let L be an upper bound on the depth of the polynomial-sized Boolean circuit which is to be evaluated. Let $\alpha \in (0,1)$ be a noise ratio with*

$$\sqrt{8(m+1)N} \leq \alpha q \leq \frac{q}{\sqrt{8}(m+1) \cdot N \cdot (N+1)^L}.$$

Then, the Dual-Regev FHE *scheme* $\Sigma = \mathsf{DualFHE}_{\mathsf{CD}}$ *with certified deletion in Construction 3 is extractable. In other words, for any efficiently computable circuit* $C : \{0,1\}^\ell \to \{0,1\}$ *and any input* $x \in \{0,1\}^\ell$*:*

$$\Pr\left[y \neq C(x_1,\dots,x_\ell) \;\middle|\; \begin{array}{c}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KeyGen}(1^\lambda,1^L)\\(\mathsf{vk},|\mathsf{CT}\rangle)\leftarrow\mathsf{Enc}(\mathsf{pk},x)\\(|\widetilde{\mathsf{CT}}\rangle,t_C)\leftarrow\mathsf{Eval}(C,|\mathsf{CT}\rangle,\mathsf{pk})\\(\varrho,y)\leftarrow\mathsf{Extract}\langle\mathcal{S}(|\widetilde{\mathsf{CT}}\rangle,t_C),\mathcal{R}(\mathsf{sk})\rangle\end{array}\right] \leq \mathrm{negl}(\lambda), \quad \textit{and}$$

$$\Pr\left[\mathsf{Vrfy}(\mathsf{vk},\pi)=\bot \;\middle|\; \begin{array}{c}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KeyGen}(1^\lambda,1^L)\\(\mathsf{vk},|\mathsf{CT}\rangle)\leftarrow\mathsf{Enc}(\mathsf{pk},x)\\(|\widetilde{\mathsf{CT}}\rangle,t_C)\leftarrow\mathsf{Eval}(C,|\mathsf{CT}\rangle,\mathsf{pk})\\(\varrho,y)\leftarrow\mathsf{Extract}\langle\mathcal{S}(|\widetilde{\mathsf{CT}}\rangle,t_C),\mathcal{R}(\mathsf{sk})\rangle\\\pi\leftarrow\mathsf{Del}(\varrho)\end{array}\right] \leq \mathrm{negl}(\lambda).$$

Proof. Let $C : \{0,1\}^\ell \to \{0,1\}$ be an efficiently computable circuit and let $x \in \{0,1\}^\ell$ be any input. Let $(\varrho, y) \leftarrow \mathsf{Extract}\langle\mathcal{S}(|\widetilde{\mathsf{CT}}\rangle, t_C), \mathcal{R}(\mathsf{sk})\rangle$ denote the outcome of the interactive protocol between the sender $\mathcal{S}$ and the receiver $\mathcal{R}$, where $(|\widetilde{\mathsf{CT}}\rangle, t_C) \leftarrow \mathsf{Eval}(C, |\mathsf{CT}\rangle, \mathsf{pk})$ is the post-evaluation state and $\mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{pk}, x)$ is the initial ciphertext for $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$. Recall that the receiver $\mathcal{R}$ reversibly performs the decryption procedure Dec (with the secret key sk hard-coded) during the execution of the protocol $\Pi = \langle\mathcal{S}(|\widetilde{\mathsf{CT}}\rangle, t_C), \mathcal{R}(\mathsf{sk})\rangle$ in Protocol 1. Therefore, it follows that the measurement outcome y is equal to $C(x_1,\dots,x_\ell)$ with overwhelming probability due Lemma 20. This shows the first property.

To show the second property, we can use the Rewinding Lemma (Lemma 19) to argue that after the interactive protocol $\Pi = \langle\mathcal{S}(\widetilde{\mathsf{CT}}, t_C), \mathcal{R}(\mathsf{sk})\rangle$ between the sender $\mathcal{S}$ and receiver $\mathcal{R}$ is complete, the sender $\mathcal{S}$ is in possession of a quantum state ϱ in system C_{in} that satisfies

$$\|\varrho - |\mathsf{CT}\rangle\langle\mathsf{CT}|\|_{\mathrm{tr}} \leq \mathrm{negl}(\lambda).$$

Therefore, the claim follows immediately from the verification correctness of Σ shown in Lemma 21. □

9.3 Proof of security

Let us now analyze the security of our FHE scheme with certified deletion in Construction 3. Note that the results in this section all essentially carry over from Section 7.2, where we analyzed the security of our Dual-Regev PKE scheme with certified deletion.

IND-CPA security of $\mathsf{DualFHE}_{\mathsf{CD}}$**.** We first prove that our scheme $\mathsf{FHE}_{\mathsf{CD}}$ in Construction 3 satisfies the notion IND-CPA security according to Definition 12. The proof is identical to the proof of IND-CPA-security of our DualPKE scheme in Theorem 6. We add it for completeness.

Theorem 9. *Let* $n \in \mathbb{N}$*, let* $q \geq 2$ *be a modulus, let* $m \geq 2n\log q$ *and let* $N = (n+1)\lceil\log q\rceil$*, each parameterized by the security parameter* $\lambda \in \mathbb{N}$*. Let* $\alpha \in (0,1)$ *be a noise ratio parameter such that* $\sqrt{8(m+1)N} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)N}}$*. Then, the scheme* $\mathsf{DualFHE}_{\mathsf{CD}}$ *in Construction 3 is* IND-CPA*-secure assuming the quantum hardness of (decisional)* $\mathsf{LWE}^{(m+1)N}_{n,q,\beta q}$*, for any* $\beta \in (0,1)$ *with* $\alpha/\beta = \lambda^{\omega(1)}$*.*

Proof. Let $\Sigma = \mathsf{DualFHE}_{\mathsf{CD}}$. We need to show that, for any QPT adversary $\mathcal{A}$, it holds that

$$\mathsf{Adv}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(0)=1] - \Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(1)=1]| \leq \mathrm{negl}(\lambda).$$

Consider the experiment $\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(b)$ between the adversary $\mathcal{A}$ and a challenger taking place as follows:

1. The challenger generates a pair $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$, and sends pk to $\mathcal{A}$.
2. $\mathcal{A}$ sends a distinct plaintext pair $(m_0, m_1) \in \{0,1\}^\ell \times \{0,1\}^\ell$ to the challenger.
3. The challenger computes $(\mathsf{vk}, \mathsf{CT}_b) \leftarrow \mathsf{DualFHE}_{\mathsf{CD}}.\mathsf{Enc}(\mathsf{pk}, m_b)$, and sends $|\mathsf{CT}_b\rangle$ to $\mathcal{A}$.
4. $\mathcal{A}$ outputs a guess $b' \in \{0,1\}$, which is also the output of the experiment.

Recall that the procedure $\mathsf{Enc}(\mathsf{pk}, m_b)$ outputs a pair $(\mathsf{vk}, |\mathsf{CT}_b\rangle)$, where

$$\left(\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}, (\mathbf{y}_1, \ldots, \mathbf{y}_N) \in \mathbb{Z}_q^{n\times N}\right) \leftarrow \mathsf{vk}$$

is the verification key and where the ciphertext $|\mathsf{CT}_b\rangle$ is within negligible trace distance of

$$\sum_{\mathbf{S}\in\mathbb{Z}_q^{n\times N}} \sum_{\mathbf{E}\in\mathbb{Z}_q^{(m+1)\times N}} \varrho_{\alpha q}(\mathbf{E})\, \omega_q^{-\mathrm{Tr}[\mathbf{S}^T\mathbf{Y}]}\, |\mathbf{A}\cdot\mathbf{S} + \mathbf{E} + m_b\cdot\mathbf{G} \pmod q\rangle\,. \tag{19}$$

Here, $\mathbf{Y} \in \mathbb{Z}_q^{n\times N}$ is the matrix composed of the rows $\mathbf{y}_1, \ldots, \mathbf{y}_N$. Let $\beta \in (0,1)$ be any parameter with $\alpha/\beta = \lambda^{\omega(1)}$. Then, it follows from Theorem 5 that, under the (decisional) $\mathsf{LWE}^{(m+1)N}_{n,q,\beta q}$ assumption, $|\mathsf{CT}_b\rangle$ is computationally indistinguishable from the state

$$\sum_{\mathbf{U}\in\mathbb{Z}_q^{(m+1)\times N}} \omega_q^{\mathrm{Tr}[\mathbf{U}^T\bar{\mathbf{X}}]}\, |\mathbf{U}\rangle\,, \quad \bar{\mathbf{X}} = (\bar{\mathbf{x}}_1, \ldots, \bar{\mathbf{x}}_N) \sim D_{\Lambda_q^{\mathbf{y}_1}(\mathbf{A}), \frac{1}{\sqrt{2}\alpha}} \times \cdots \times D_{\Lambda_q^{\mathbf{y}_N}(\mathbf{A}), \frac{1}{\sqrt{2}\alpha}}. \tag{20}$$

Here $(\bar{\mathbf{x}}_1, \ldots, \bar{\mathbf{x}}_N)$ refer to the rows of the matrix $\bar{\mathbf{X}} \in \mathbb{Z}_q^{(m+1)\times N}$. Finally, because the state in Eq. (20) is completely independent of the bit $b \in \{0,1\}$, it follows that

$$\mathsf{Adv}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

This proves the claim. □

IND-CPA-CD security of $\mathsf{DualFHE}_{\mathsf{CD}}$. Let us now analyze the security of our Dual-Regev homomorphic encryption scheme $\mathsf{DualFHE}_{\mathsf{CD}}$ in Construction 3. We prove that it satisfies *certified deletion security* assuming the *Strong Gaussian-Collapsing (SGC) Conjecture* (see Conjecture 5.2). This is a strengthening of the Gaussian-collapsing property which we proved under the (decisional) LWE assumption (see Theorem 4). The proof is similar to the proof of Theorem 7. We add it for completeness.

Theorem 10. *Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n \in \mathbb{N}$, let $q \geq 2$ be a prime modulus and $m \geq 2n\log q$. Let $N = (n+1)\cdot\lceil\log q\rceil$ be an integer and let L be an upper bound on the depth of the polynomial-sized Boolean circuit which is to be evaluated. Let $\alpha \in (0,1)$ be a noise ratio such that*

$$\sqrt{8(m+1)N} \leq \alpha q \leq \frac{q}{\sqrt{8}(m+1)\cdot N\cdot(N+1)^L}.$$

Then, the Dual-Regev homomorphic encryption scheme $\mathsf{DualFHE}_{\mathsf{CD}}$ in Construction 3 is IND-CPA-CD-secure assuming the Strong Gaussian-Collapsing property $\mathsf{SGC}^N_{n,(m+1),q,\frac{1}{\alpha}}$ from Conjecture 5.2.

Proof. Let $\Sigma = \mathsf{DualFHE}_{\mathsf{CD}}$. We need to show that, for any QPT adversary $\mathcal{A}$, it holds that

$$\mathsf{Adv}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{Exp}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{Exp}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

We consider the following sequence of hybrids:

$\mathbf{H_0}$: This is the experiment $\mathsf{Exp}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(0)$ between $\mathcal{A}$ and a challenger:

1. The challenger samples a random matrix $\bar{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and a vector $\bar{\mathbf{x}} \xleftarrow{\$} \{0,1\}^m$ and chooses $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}}\cdot\bar{\mathbf{x}} \pmod q]^T$. The challenger chooses the secret key $\mathsf{sk} \leftarrow (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$ and the public key $\mathsf{pk} \leftarrow \mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}$.
2. $\mathcal{A}$ sends a distinct plaintext pair $(m_0, m_1) \in \{0,1\}\times\{0,1\}$ to the challenger. (Note: Without loss of generality, we can just assume that $m_0 = 0$ and $m_1 = 1$).
3. The challenger runs $(|\psi_{\mathbf{y}_i}\rangle, \mathbf{y}_i) \leftarrow \mathsf{GenPrimal}(\mathbf{A}^T, \sigma)$ in Algorithm 2, for $i \in [N]$, and outputs
$$\left(\mathsf{vk} \leftarrow (\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}, (\mathbf{y}_1, \ldots, \mathbf{y}_N) \in \mathbb{Z}_q^{n\times N}), \quad |\mathsf{CT}_0\rangle \leftarrow |\psi_{\mathbf{y}_1}\rangle \otimes \cdots \otimes |\psi_{\mathbf{y}_N}\rangle\right).$$
4. At some point in time, $\mathcal{A}$ returns a certificate $\pi = (\pi_1, \ldots, \pi_N)$ to the challenger.
5. The challenger outputs $\top$, if $\mathbf{A}^T \cdot \pi_i = \mathbf{y}_i \pmod q$ and $\|\pi_i\| \leq \sqrt{m+1}/\sqrt{2}\alpha$ for $i \in [N]$, and outputs $\bot$, otherwise. If π passes the test with outcome $\top$, the challenger sends sk to $\mathcal{A}$.
6. $\mathcal{A}$ outputs a guess $b' \in \{0,1\}$, which is also the output of the experiment.

$\mathbf{H_1}$: This is same experiment as in $\mathbf{H_0}$, except that (in Step 3) the challenger prepares the ciphertext in the Fourier basis rather than the standard basis. In other words, $\mathcal{A}$ receives the pair

$$\left(\mathsf{vk} \leftarrow (\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}, (\mathbf{y}_1, \ldots, \mathbf{y}_N) \in \mathbb{Z}_q^{n\times N}), \quad |\mathsf{CT}_0\rangle \leftarrow \mathsf{FT}_q|\psi_{\mathbf{y}_1}\rangle \otimes \cdots \otimes \mathsf{FT}_q|\psi_{\mathbf{y}_N}\rangle\right).$$

$\mathbf{H_2}$: This experiment is an N-fold variant of $\mathsf{StrongGaussCollapseExp}_{\mathcal{H},\mathcal{D},\lambda}(0)$ in Conjecture 5.2:

1. The challenger samples a random matrix $\bar{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and a vector $\bar{\mathbf{x}} \xleftarrow{\$} \{0,1\}^m$ and chooses $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}}\cdot\bar{\mathbf{x}} \pmod q]$ and $\mathbf{t} = (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$.
2. The challenger runs $(|\hat{\psi}_{\mathbf{y}_i}\rangle, \mathbf{y}_i) \leftarrow \mathsf{GenDual}(\mathbf{A}^T, \sigma)$ in Algorithm 1, for $i \in [N]$, and sends the following tiplet to the adversary $\mathcal{A}$:
$$\left(|\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes |\hat{\psi}_{\mathbf{y}_N}\rangle, \quad \mathbf{A}^T \in \mathbb{Z}_q^{n\times(m+1)}, \quad \mathbf{Y} = (\mathbf{y}_1, \ldots, \mathbf{y}_N) \in \mathbb{Z}_q^{n\times N}\right).$$
3. At some point in time, $\mathcal{A}$ returns a certificate π to the challenger.
4. The challenger outputs $\top$, if $\mathbf{A}^T \cdot \pi_i = \mathbf{y}_i \pmod q$ and $\|\pi_i\| \leq \sqrt{m+1}/\sqrt{2}\alpha$ for $i \in [N]$, and outputs $\bot$, otherwise. If π passes the test with outcome $\top$, the challenger sends sk to $\mathcal{A}$.
5. $\mathcal{A}$ outputs a guess $b' \in \{0,1\}$, which is also the output of the experiment.

$\mathbf{H_3}$: This is an N-fold variant of the experiment in $\mathsf{StrongGaussCollapseExp}_{\mathcal{H},\mathcal{D},\lambda}(1)$ in Conjecture 5.2; it is the same as $\mathbf{H_2}$, except that the states $|\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes |\hat{\psi}_{\mathbf{y}_N}\rangle$ (in Step 2) are measured in the computational basis before they are sent to $\mathcal{A}$.

$\mathbf{H_4}$: This is same experiment as $\mathbf{H_3}$, except that (in Step 2) the challenger additionally applies the Pauli operators $\mathbf{Z}_q^{\mathbf{g}_1} \otimes \cdots \otimes \mathbf{Z}_q^{\mathbf{g}_N}$ to the states $|\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes |\hat{\psi}_{\mathbf{y}_N}\rangle$ before they are measured in the computational basis, where $(\mathbf{g}_1, \ldots, \mathbf{g}_N)$ are the rows of the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ in Eq. (18).

$\mathbf{H_5}$: This is same experiment as $\mathbf{H_4}$, except that (in Step 2) $\mathcal{A}$ receives the triplet

$$\left(\mathbf{Z}_q^{\mathbf{g}_1} |\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes \mathbf{Z}_q^{\mathbf{g}_N} |\hat{\psi}_{\mathbf{y}_N}\rangle, \quad \mathbf{A}^T \in \mathbb{Z}_q^{n\times(m+1)}, \quad \mathbf{y} = (\mathbf{y}_1, \ldots, \mathbf{y}_N) \in \mathbb{Z}_q^{n\times N}\right).$$

$\mathbf{H_6}$: This is same experiment as $\mathbf{H_5}$, except that (in Step 2) the challenger prepares the quantum states in the Fourier basis instead. In other words, $\mathcal{A}$ receives the triplet

$$\left(\mathsf{FT}_q^\dagger \mathbf{Z}_q^{\mathbf{g}_1} |\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes \mathsf{FT}_q^\dagger \mathbf{Z}_q^{\mathbf{g}_N} |\hat{\psi}_{\mathbf{y}_N}\rangle, \quad \mathbf{A}^T \in \mathbb{Z}_q^{n\times(m+1)}, \quad \mathbf{y} = (\mathbf{y}_1, \ldots, \mathbf{y}_N) \in \mathbb{Z}_q^{n\times N}\right).$$

$\mathbf{H_7}$: This is the experiment $\mathsf{Exp}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1)$.

We now show that the hybrids are indistinguishable.

Claim 9.

$$\Pr[\mathsf{Exp}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(0) = 1] = \Pr[\mathbf{H_1} = 1].$$

Proof. Without loss of generality, we can assume that $\mathcal{A}$ applies the inverse Fourier transform immediately upon receiving the quantum ciphertext. Therefore, the success probabilities are identical in $\mathbf{H_0}$ and $\mathbf{H_1}$. □

Claim 10.

$$\Pr[\mathbf{H_1} = 1] = \Pr[\mathbf{H_2} = 1].$$

Proof. Because the challenger in $\mathbf{H_1}$ always sends the ciphertext $|\mathsf{CT}_0\rangle$ corresponding to $m_0 = 0$ to the adversary $\mathcal{A}$, the two hybrids $\mathbf{H_1}$ and $\mathbf{H_2}$ are identical. □

Claim 11. *Under the Strong Gaussian-Collapsing property* $\mathsf{SGC}^N_{n,(m+1),q,\frac{1}{\alpha}}$, *it holds that*

$$|\Pr[\mathbf{H_2} = 1] - \Pr[\mathbf{H_3} = 1]| \leq \mathrm{negl}(\lambda).$$

Proof. This follows from Conjecture 5.2. □

Claim 12.

$$\Pr[\mathbf{H_3} = 1] = \Pr[\mathbf{H_4} = 1].$$

Proof. Because the challenger measures the state $|\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes |\hat{\psi}_{\mathbf{y}_N}\rangle$ in Step 2 in the computational basis, applying the phase operators $\mathbf{Z}_q^{\mathbf{g}_1} \otimes \cdots \otimes \mathbf{Z}_q^{\mathbf{g}_N}$ before the measurement does not affect the outcome. □

Claim 13. *Under the Strong Gaussian-Collapsing property* $\mathsf{SGC}^N_{n,(m+1),q,\frac{1}{\alpha}}$, *it holds that*

$$|\Pr[\mathbf{H_4} = 1] - \Pr[\mathbf{H_5} = 1]| \leq \mathrm{negl}(\lambda).$$

Proof. This follows from Conjecture 5.2 since, without loss of generality, we can assume that $\mathcal{A}$ applies the phase operators $\mathbf{Z}_q^{\mathbf{g}_1} \otimes \cdots \otimes \mathbf{Z}_q^{\mathbf{g}_N}$ immediately upon receiving the states $|\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes |\hat{\psi}_{\mathbf{y}_N}\rangle$ as input. □

Claim 14.

$$\Pr[\mathbf{H_5} = 1] = \Pr[\mathbf{H_6} = 1].$$

Proof. Without loss of generality, we can assume that $\mathcal{A}$ applies the Fourier transform immediately upon receiving $\mathbf{Z}_q^{\mathbf{g}_1} |\hat{\psi}_{\mathbf{y}_1}\rangle \otimes \cdots \otimes \mathbf{Z}_q^{\mathbf{g}_N} |\hat{\psi}_{\mathbf{y}_N}\rangle$. Therefore, the success probabilities in $\mathbf{H_5}$ and $\mathbf{H_6}$ are identical. □

Claim 15.

$$|\Pr[\mathbf{H_6} = 1] - \Pr[\mathsf{Exp}^{\mathsf{pk\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathrm{negl}(\lambda).$$

Proof. From Lemma 6, we have $\mathsf{FT}_q \mathbf{X}_q^{\mathbf{v}} = \mathbf{Z}_q^{\mathbf{v}} \mathsf{FT}_q$, for all $\mathbf{v} \in \mathbb{Z}_q^m$. Hence, in $\mathbf{H_6}$, we can instead assume that the challenger runs $(|\psi_{\mathbf{y}_i}\rangle, \mathbf{y}_i) \leftarrow \mathsf{GenPrimal}(\mathbf{A}^T, 1/\alpha)$ in Algorithm 2, for $i \in [N]$, and then sends the following to $\mathcal{A}$:

$$\left(\mathsf{vk} \leftarrow (\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}, (\mathbf{y}_1, \ldots, \mathbf{y}_N) \in \mathbb{Z}_q^{n\times N}), \quad |\mathsf{CT}_1\rangle \leftarrow \mathbf{X}_q^{\mathbf{g}_1} |\psi_{\mathbf{y}_1}\rangle \otimes \cdots \otimes \mathbf{X}_q^{\mathbf{g}_N} |\psi_{\mathbf{y}_N}\rangle\right).$$

From Corollary 1, it follows that the states $\mathsf{FT}_q^\dagger \mathbf{Z}_q^{\mathbf{v}} |\hat{\psi}_{\mathbf{y}}\rangle$ and $\mathbf{X}_q^{\mathbf{v}} |\psi_{\mathbf{y}}\rangle$ are within negligible trace distance, for all $\mathbf{v} \in \mathbb{Z}_q^m$. Because the challenger in $\mathbf{H_7}$ always sends $|\mathsf{CT}_1\rangle$ corresponding to $m_1 = 1$ to the adversary $\mathcal{A}$, it follows that the distinguishing advantage between $\mathbf{H_6}$ and $\mathbf{H_7} = \mathsf{Exp}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A},\lambda}(1)$ is negligible. □

Because the hybrids $\mathbf{H_0}$ and $\mathbf{H_7}$ are indistinguishable, this implies that

$$\mathsf{Adv}^{\mathsf{he\text{-}cert\text{-}del}}_{\Sigma,\mathcal{A}}(\lambda) \leq \mathrm{negl}(\lambda).$$

□

References

[Aar16] Scott Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, 2016.

[AC12] Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces, 2012.

[AJOP20] Gorjan Alagic, Stacey Jeffery, Maris Ozols, and Alexander Poremba. On quantum chosen-ciphertext attacks and learning with errors. *Cryptography*, 4(1), 2020.

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 99–108. ACM, 1996.

[AP20] Prabhanjan Ananth and Rolando L. La Placa. Secure software leasing, 2020.

[Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–636, 1993.

[BB84] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, page 175, India, 1984.

[BCM+21] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device, 2021.

[BI20] Anne Broadbent and Rabib Islam. Quantum encryption with certified deletion. *Lecture Notes in Computer Science*, page 92–122, 2020.

[BPTG14] Raphael Bost, Raluca Ada Popa, Stephen Tu, and Shafi Goldwasser. Machine learning classification over encrypted data. *IACR Cryptology ePrint Archive*, 2014:331, 2014.

[Bra18] Zvika Brakerski. Quantum fhe (almost) as secure as classical. Cryptology ePrint Archive, Report 2018/338, 2018. https://ia.cr/2018/338.

[BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. In *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, FOCS '11, page 97–106, USA, 2011. IEEE Computer Society.

[CFGN96] Ran Canetti, Uri Feige, Oded Goldreich, and Moni Naor. Adaptively secure multi-party computation. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 639–648, New York, NY, USA, 1996. Association for Computing Machinery.

[CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography, 2021.

[CLZ21] Yilei Chen, Qipeng Liu, and Mark Zhandry. Quantum algorithms for variants of average-case lattice problems via filtering, 2021.

[CMP20] Andrea Coladangelo, Christian Majenz, and Alexander Poremba. Quantum copy-protection of compute-and-compare programs in the quantum random oracle model, 2020.

[CRW19] Xavier Coiteux-Roy and Stefan Wolf. Proving erasure. *2019 IEEE International Symposium on Information Theory (ISIT)*, Jul 2019.

[DKW11] Stefan Dziembowski, Tomasz Kazana, and Daniel Wichs. One-time computable self-erasing functions. In *Theory of Cryptography - 8th Theory of Cryptography Conference, TCC 2011*, volume 6597 of *Lecture Notes in Computer Science*, page 125. Springer, 2011.

[FM18] Honghao Fu and Carl A. Miller. Local randomness: Examples and application. *Physical Review A*, 97(3), Mar 2018.

[Gen09] Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. crypto.stanford.edu/craig.

[GGV20] Sanjam Garg, Shafi Goldwasser, and Prashant Nalini Vasudevan. Formalizing data deletion in the context of the right to be forgotten. *IACR Cryptol. ePrint Arch.*, page 254, 2020.

[GKZ19] Alex B. Grilo, Iordanis Kerenidis, and Timo Zijlstra. Learning-with-errors problem is easy with quantum samples. *Physical Review A*, 99(3), Mar 2019.

[GMP22] Alexandru Gheorghiu, Tony Metger, and Alexander Poremba. Quantum cryptography with classical communication: parallel remote state preparation for copy-protection, verification, and more, 2022.

[GPV07] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. Cryptology ePrint Archive, Report 2007/432, 2007. https://eprint.iacr.org/2007/432.

[GR02] Lov K. Grover and Terry Rudolph. Creating superpositions that correspond to efficiently integrable probability distributions. *arXiv: Quantum Physics*, 2002.

[GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. Cryptology ePrint Archive, Report 2013/340, 2013. https://ia.cr/2013/340.

[HH00] L. Hales and S. Hallgren. An improved quantum fourier transform algorithm and applications. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 515–525, 2000.

[HILL88] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions. In *PROC. 20TH STOC*, pages 12–24, 1988.

[HMNY21a] Taiga Hiroka, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Certified everlasting zero-knowledge proof for qma, 2021.

[HMNY21b] Taiga Hiroka, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Quantum encryption with certified deletion, revisited: Public key, attribute-based, and classical communication, 2021.

[JL00] Stanisław Jarecki and Anna Lysyanskaya. Adaptively secure threshold cryptography: Introducing concurrency, removing erasures. In *Proceedings of the 19th International Conference on Theory and Application of Cryptographic Techniques*, EUROCRYPT'00, page 221–242, Berlin, Heidelberg, 2000. Springer-Verlag.

[KNY21] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Secure software leasing from standard assumptions, 2021.

[KRS09] Robert Konig, Renato Renner, and Christian Schaffner. The operational meaning of min- and max-entropy. *IEEE Transactions on Information Theory*, 55(9):4337–4347, sep 2009.

[LZ19] Qipeng Liu and Mark Zhandry. Revisiting post-quantum fiat-shamir. Cryptology ePrint Archive, Paper 2019/262, 2019. https://eprint.iacr.org/2019/262.

[Mah18] Urmila Mahadev. Classical verification of quantum computations, 2018.

[MQU07] Jörn Müller-Quade and Dominique Unruh. Long-term security and universal composability. In Salil P. Vadhan, editor, *Theory of Cryptography*, pages 41–60, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.

[MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 372–381, 2004.

[MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.

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

[PT10] Daniele Perito and Gene Tsudik. Secure code update for embedded devices via proofs of secure erasure. Cryptology ePrint Archive, Report 2010/217, 2010. https://ia.cr/2010/217.

[RAD78] R L Rivest, L Adleman, and M L Dertouzos. On data banks and privacy homomorphisms. *Foundations of Secure Computation, Academia Press*, pages 169–179, 1978.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):34:1–34:40, 2005.

[Rob19] Bhaskar Roberts. Toward secure quantum money. Princeton University Senior Thesis, 2019. http://arks.princeton.edu/ark:/88435/dsp01nc580q51r.

[SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. Cryptology ePrint Archive, Paper 2009/285, 2009. https://eprint.iacr.org/2009/285.

[TL17] Marco Tomamichel and Anthony Leverrier. A largely self-contained and complete security proof for quantum key distribution. *Quantum*, 1:14, July 2017.

[Unr13] Dominique Unruh. Revocable quantum timed-release encryption. Cryptology ePrint Archive, Report 2013/606, 2013. https://ia.cr/2013/606.

[Unr15] Dominique Unruh. Computationally binding quantum commitments. Cryptology ePrint Archive, Paper 2015/361, 2015. https://eprint.iacr.org/2015/361.

[Wat06] John Watrous. Zero-knowledge against quantum attacks. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, page 296–305, New York, NY, USA, 2006. Association for Computing Machinery.

[Wil13] Mark M. Wilde. *Quantum Information Theory*. Cambridge University Press, USA, 1st edition, 2013.

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