Revocable Cryptography in a Quantum World

Thesis by

Alexander Mario Poremba

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy in Computer Science

Caltech

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2023
Defended May 16, 2023

ACKNOWLEDGEMENTS

This thesis would not be possible without the endless support from many friends and colleagues.

Most importantly, I want to thank my advisor, Thomas Vidick, for his wonderful supervision and support over the years. The journey from incoming graduate student to independent researcher is long and difficult at times, and I am tremendously grateful for his guidance along the way. Thomas gave me unlimited freedom to pursue whatever research directions I found interesting; gave me the space to work on them—always checking in on me to make sure I was on the right path; provided detailed feedback on countless drafts I have sent him over the years, and always showed concern for my well-being, which made it easy to pull through even difficult times. Thomas is an exceptional advisor; he taught me how to be critical with my own work; urged me to travel during the summers to seek out external collaborations, and helped organize many research visits. I was very fortunate to have him as my advisor and I will always admire his curiosity and breadth of expertise.

I also want to thank the other members of my thesis committee: Urmila Mahadev, John Preskill, and Chris Umans. My research has benefited a lot from their teaching and countless interactions with them over the years. I am very grateful for having shared the same department with them.

During my PhD, I was fortunate to collaborate with many incredibly talented people who taught me a lot (in no particular order): Sumeet Khatri, Ryan LaRose, Lukasz Cincio, Andrew Sornborger, Patrick Coles, Gorjan Alagic, Stacey Jeffery, Maris Ozols, Marco Cerezo, Andrea Coladangelo, Christian Majenz, Alexandru Gheorghiu, Tony Metger, András Gilyén, Prabhanjan Ananth, Vinod Vaikuntanathan, James Bartusek, Dakshita Khurana, Giulio Malavolta, Michael Walter, Chen Bai, Kaiyan Shi, John Bostanci, Yuval Efron, Luowen Qian, and Henry Yuen.

I want to thank all the quantum information students and post-docs at Caltech for many fun and illuminating conversations (in no particular order): Andrea Coladangelo, Alexandru Gheorgiu, Richard Kueng, Anand Natarajan, Tina Zhang, András Gilyén, Hsin-Yuan (Robert) Huang, Alex Jahn, Robbie King, Joe Slote, Jiaqing Jiang, Atul Singh Arora, Ulysse Chabaud, Aleksander Kubica, Jenish Mehta, Spencer Gordon, Jiayu Zhang, Laura Lewis, Victor Albert, Eugene Tang, Joseph Iverson, Alex Buser, Shouzhen (Bailey) Gu, Chris Pattison, and Chi-Fang (Anthony) Chen.

I want to thank Prabhanjan Ananth, Dakshita Khurana, Vinod Vaikuntanathan, and Henry Yuen for hosting me during (often multiple) research visits. My research has benefited tremendously from my collaborations with them over the years. I also want to thank them for sharing their perspectives about research, the future and life in general.

I want to thank Los Alamos National Laboratory for hosting me during the summer before my PhD. I was very fortunate to have been part of the first quantum computing summer school. Thanks to my advisors Patrick Coles, Lukasz Cincio, and Andrew Sornborger, and all my wonderful colleagues for an unforgettable and productive summer: Adam Pearson, Arkin Tikku, Christina Cirstoiu, Joan Étude Arrow, Joe Iosue, Namit Anand, Ryan LaRose, Sumeet Khatri and Yousif Almulla.

I want to thank the Simons Institute for the Theory of Computing for hosting me multiple times. Many of my collaborations started out as random interactions at the institute, particularly during the 2020 "Quantum Wave in Computing" spring program as well as the 2022 summer program "Extended Reunion: The Quantum Wave in Computing."

I had an extremely happy time throughout my PhD at Caltech, and for that I have many people in my cohort to thank for: Daniel Leibovici, Nikola Kovachki, Musab Jilani, Dmitry Burov, Christoph Bauinger, Matt Levine, Yousuf Soliman, Jing Shuang (Lisa) Li, and Apurva Badithela.

Throughout my time at Caltech I received a lot of help from many incredible administrators at CMS and the ISP office. Special thanks to Maria Lopez, Bonnie Leung, Laura Flower Kim and Daniel Yoder for all the hard work and support—Caltech is lucky to have them! My sincere gratitude also goes to Alice Sogomonian at the Student Health Services, as well as Jacob Dalton, Natalie Gilmore, and Kevin Lee who administer the Caltech emergency fund.

Finally, I would like to thank my parents and my brother for their wonderful support over the years. More than anybody, they have kept me sane and motivated to pull through my PhD.

ABSTRACT

Quantum cryptography leverages unique features of quantum mechanics in order to construct cryptographic primitives which are oftentimes impossible for digital computers. Cryptographic applications of quantum computers therefore have the potential for useful quantum advantage—entirely without computational speed-ups. Can we use the power of quantum states to address fundamental limitations in the world of classical cryptography, such as the intricate problem of "revoking" information from an untrusted party? This thesis undertakes a systematic study of how to delegate and revoke privileges in a world in which quantum computers become widely available. As part of a single framework we call *revocable cryptography*, we show how to revoke programs, encrypted data, and even cryptographic keys under standard assumptions.

In the first part of this thesis, we focus on the following question: can we use the no-cloning principle of quantum mechanics and encode a program in such a way that it can be evaluated, yet it cannot be *pirated*? Naturally, we would also like to ensure that, once the program is "returned," the recipient loses its ability to evaluate it. While this quantum notion of *secure software leasing* (SSL) was shown to be impossible for general programs by Ananth and La Placa (Eurocrypt 2021), their work left open the possibility that it is achievable for more primitive classes of programs. We construct an SSL scheme for a large class of evasive functions known as *compute-and-compare programs*—a more expressive generalization of point functions. Our scheme can be instantiated with any cryptographic hash function, and we prove its security in the quantum random oracle model. As a complementary result, we also construct a *quantum copy-protection* scheme for multi-bit point functions, which achieves a related but stronger notion of software protection previously introduced by Aaronson (CCC 2009).

In the second part of this thesis, we ask: is it possible to provably delete information by leveraging the laws of quantum mechanics? We revisit a cryptographic notion called *certified deletion*, which was proposed by Broadbent and Islam (TCC 2020). While this remarkable notion allows a classical verifier to be convinced that quantum ciphertext has been deleted by an untrusted party, it offers no additional layer of functionality. We use Gaussian superpositions over lattices to construct the first fully homomorphic encryption scheme with certified deletion – a protocol which allows an untrusted quantum server to compute on encrypted data and to also prove data deletion to a client. Our scheme has the desirable property that verification of a deletion certificate is *public*; meaning anyone can verify whether deletion has taken place. Assuming the quantum subexponential hardness of the learning with errors problem (Regev, STOC 2005), we can prove that our scheme achieves a particularly strong *information-theoretic* deletion guarantee; namely, once a valid deletion certificate is presented, the plaintext remains hidden even if the adversary is subsequently allowed to run in unbounded time.

In the final part of this thesis, we ask: is it possible to revoke a crytographic key by using the power of quantum information? We give an affirmative answer to this question and design cryptosystems with key-revocation capabilities; specifically, we consider schemes with the guarantee that, once the secret key (represented as a quantum state) is successfully revoked from a user, they no longer have the ability to perform the same functionality as before. We define and construct several fundamental cryptographic primitives with key-revocation capabilities, namely pseudorandom functions, secret-key and public-key encryption, and even fully homomorphic encryption, assuming the subexponential hardness of the learning with errors problem. Central to all our constructions is our approach for making the Dual-Regev encryption scheme (Gentry, Peikert and Vaikuntanathan, STOC 2008) revocable.

PUBLISHED CONTENT AND CONTRIBUTIONS

The chapters of this thesis are based on the following publications or preprints.

Chapter 3: This is based on the following work:

- Andrea Coladangelo, Christian Majenz, and Alexander Poremba. "Quantum copy-protection of compute-and-compare programs in the quantum random oracle model." In: *arXiv preprint arXiv:2009.13865* (2020).

  This work is based on a collaboration which started at the "Quantum Wave in Computing" 2020 spring program at the Simons Institute for the Theory of Computing. It was presented as a contributed talk at QIP 2021. Each author made significant contributions to the paper.

Chapter 4: This is based on the following works:

- Alexander Poremba. "Quantum Proofs of Deletion for Learning with Errors." In: *Proceedings of the 14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA. Ed. by Yael Tauman Kalai. Vol. 251. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 90:1-90:14.*

  This work was presented at QIP 2022, QCRYPT 2022, the ASIACRYPT Quantum Cryptography Workshop 2022, and ITCS 2023.

- James Bartusek, Dakshita Khurana, and Alexander Poremba. "Publicly-Verifiable Deletion via Target-Collapsing Functions." In: *arXiv preprint arXiv:2303.08676* (2023).

  This is a collaboration with James Bartusek and Dakshita Khurana, and came about after a visit to the University of Illinois, Urbana-Champaign. In a joint effort, we developed new techniques which allowed us to prove the "strong Gaussian-collapsing conjecture," a property of the Ajtai hash function which I posed in my paper above. To appear in CRYPTO 2023.

Chapter 5: This is based on the following work:

- Prabhanjan Ananth, Alexander Poremba, and Vinod Vaikuntanathan. "Revocable Cryptography from Learning with Errors." In: *arXiv preprint arXiv:2302.14860* (2023).

  This is a collaboration which started at the 2022 summer program "Extended Reunion: The Quantum Wave in Computing" at the Simons Institute for the Theory of Computing. Much of the work was completed in a joint effort during my visits to UC Santa Barbara, where I visted Prabhanjan Ananth, and MIT, where I visited Vinod Vaikuntanathan.

The following works were completed during my PhD, but are not included in this thesis:

[1] James Bartusek et al. *Weakening Assumptions for Publicly-Verifiable Deletion*. 2023. arXiv: 2304.09846 [quant-ph].

[2] Alexandru Gheorghiu, Tony Metger, and Alexander Poremba. *Quantum cryptography with classical communication: parallel remote state preparation for copy-protection, verification, and more*. 2022. arXiv: 2201.13445 [quant-ph].

[3] András Gilyén and Alexander Poremba. *Improved Quantum Algorithms for Fidelity Estimation*. 2022. arXiv: 2203.15993 [quant-ph].

[4] Gorjan Alagic et al. "On Quantum Chosen-Ciphertext Attacks and Learning with Errors". In: *Cryptography* 4.1 (2020). ISSN: 2410-387X. DOI: 10.3390/cryptography4010010. URL: https://www.mdpi.com/2410-387X/4/1/10.

[5] Marco Cerezo et al. "Variational Quantum Fidelity Estimation". In: *Quantum* 4 (Mar. 2020), p. 248. ISSN: 2521-327X. DOI: 10.22331/q-2020-03-26-248. URL: https://doi.org/10.22331/q-2020-03-26-248.

[6] Sumeet Khatri et al. "Quantum-assisted quantum compiling". In: *Quantum* 3 (May 2019), p. 140. ISSN: 2521-327X. DOI: 10.22331/q-2019-05-13-140. URL: https://doi.org/10.22331/q-2019-05-13-140.

TABLE OF CONTENTS

Chapter 1

INTRODUCTION

Quantum computers have the potential to completely transform disciplines such as physics and material science [47, 95]. While large-scale quantum computing is at least several decades away, the promise of quantum advantage has spurred an interest in developing quantum algorithms for problems in quantum chemistry, optimization, and machine learning. To this day, however, it is still unclear whether these areas admit exponential speed-ups for practically relevant problems [1, 94]. Moreover, due to the massive overhead required for quantum error-correction, it seems that mere polynomial speed-ups will not be of practical relevance for perhaps several decades to come [124].

Quantum cryptography has also received significant attention by both industry and academia in the advent of quantum key distribution (QKD) [29], which enables secure communication from the laws of physics alone. This is in contrast with conventional public-key encryption systems that require computational assumptions, such as the hardness of factoring, discrete log or worst-case lattice problems. While QKD is frequently criticised [118] for being less practical compared to classical post-quantum alternatives, recent advances in quantum cryptography have also given rise to entirely new primitives which have no classical counterpart. At the heart of these new primitives lies the *no-cloning principle* of quantum mechanics [135, 56] which stipulates that it is fundamentally impossible to copy an unknown quantum state. In his seminal work from the 1970s, Wiesner [132] proposed a quantum money scheme, wherein quantum states are used to construct banknotes that can be verified but cannot be counterfeited. Ever since this watershed moment, and especially so in recent years, a wide variety of so-called *unclonable* primitives [2, 4, 128, 27, 41, 40, 17, 70] have been studied and constructed. Due to its quantum nature, *unclonable cryptography* seems to offer an alternative path towards quantum advantage—entirely without computational speed-ups.

Following Wiesner's work, Aaronson [2] introduced the idea of *public-key* quantum money as a means of generating unforgeable quantum banknotes that anyone (not just the bank) can verify. In the very same work, Aaronson also proposed the idea of *quantum copy-protection* to prevent software piracy. Ananth and La Placa [17] later introduced a weaker form of software protection called *secure software leasing* in which the quantum program is eventually returned and verified. Another line of work [75, 128, 40, 41] exploits no-cloning to protect ciphertexts from being replicated. Gottesman [75] proposed the notion of *unclonable encryption*—a concept that was recently revisited by Broadbent and Lord [41]. Unruh [128] gave a quantum timed-release encryption scheme that is *revocable*: it enables a user to return a timed-release encryption before some fixed amount of time has passed, thereby losing all access to the data. Broadbent and Islam [40] constructed a quantum encryption scheme with *certified deletion* which is inspired by the QKD protocol [29, 122]. This cryptographic notion is made possible by the principle of complementarity of quantum mechanics, which ensures that one cannot measure two mutually incompatible observables at the same time. Contrary to Unruh's [127] notion of revocable quantum ciphertexts, certificates of deletion are entirely classical. The security definition requires that, once a valid certificate is presented, the plaintext remains hidden even if the secret key is later revealed. Hiroka et al. [83] later constructed more advanced primitives with certified deletion; namely, public-key and attribute-based encryption. In later work, Bartusek and Khurana [23] considered generic transformations for encryption schemes with certified deletion. Building on the work of Broadbent and Islam [41], they use Wiesner's conjugate coding to construct advanced encryption systems with the stronger notion of *certified everlasting security*, which allows the adversary to be unbounded once deletion is successful. A similar notion of everlasting security was previously considered by Hiroka et al. [82] who studied certified everlasting zero-knowledge proofs for QMA.

A recent series of works studied unclonable primitives in the context of advanced functionalities, such as digital signatures, decryption and pseudorandom functions. Ben-David and Sattath [27] proposed *quantum signature tokens* that prevent a recipient from signing more than one message at a time. Georgiou and Zhandry [70] considered unclonable decryption keys. Coladangelo et al. [53] constructed a copy-protection scheme for pseudorandom functions using *subspace coset states*, which can be seen as an extension of Wiesner's conjugate coding technique.

Unlike in classical cryptography, where many fundamental (and even advanced) cryptographic primitives can be based solely on the hardness of lattice problems, primarily in the form of the *learning with errors* assumption [112], the situation is quite different in the world of unclonable cryptography. While some unclonable primitives are achievable information-theoretically in restricted settings, such as in the private-key setting [41, 40, 70, 23, 88] or with respect to weaker notions of software protection [52, 89, 42], most advanced primitives either require strong cryptographic assumptions [70, 54, 119], or rely on unproven conjectures [4, 136]. This is especially the case for unclonable primitives with strong *functionalities*, such as public-key quantum money [4, 136, 119], copy-protection of pseudorandom functions and digital signatures [54, 96] or unclonable decryption keys [70, 54] which require non-standard assumptions such as indistinguishability obfuscation [22] or extractable witness encryption [64]. On the contrary, unclonbable primitives with more limited functionalities, such as weaker variants of public-key quantum money [110, 116], unclonable public-key encryption schemes [14] or public-key encryption schemes with certified deletion [23, 84] are achievable under standard assumptions, such as in the quantum random oracle model [34] or from the learning with errors assumption [112].

To this day, many unclonable primitives are still either directly [40, 41, 70] or indirectly [90, 82, 53] rooted in Wiesner's conjugate coding technique, and can only achieve advanced functionalities when combined with strong building blocks from the world of classical cryptography. In particular, constructing advanced primitives such as public-key quantum money or copy-protection schemes from lattices remains a fundamental open problem in cryptography. This raises the following questions: Can we design advanced unclonable primitives using techniques that go beyond Wiesner's conjugate coding approach? If so, is it possible to base these primitives on standard assumptions, such as the hardness of worst-case lattice problems? And lastly, is it possible to use quantum information to introduce new features to cryptographic primitives, beyond unclonability?

This thesis

In the past few decades, we have witnessed the birth of remarkable cryptographic primitives, such as secure multi-party computation [28], zero-knowledge proof systems [74, 98], and even fully homomorphic encryption [113, 66, 37]. However, despite a lot progress, several fundamental problems still seem remain out of reach for classical cryptography. For instance, can we use cryptography to prevent software piracy—a problem that accounts for billions of dollars of losses every year? Can we revoke decryption priviliges from a network of users? Can we certify that user data stored on a remote cloud server has been deleted? Data protection, in particular, has become a major challenge in today's age of cloud computing and artificial intelligence. Collectively, all of these problems amount to a single fundamental question, namely: how can we "revoke" sensitive information from an untrusted party? If the information at hand is represented in terms of classical bits, then such a task is clearly impossible to achieve on conventional digital computers.

The central goal of this thesis is to understand how to delegate and revoke privileges in a world in which quantum computers become widely available. We make progress on the following questions, in particular: Can we use the power of quantum states in order to encode useful information that can later be revoked? How can we go about formalizing the notion that certain privileges have been revoked? Is it possible to provide meaningful guarantees for revocation, particularly in the context of programs and decryption keys which offer additional functionalities? We show that quantum computers are uniquely capable at addressing all of these questions.

The framework: Revocable cryptography

This thesis continues a recent line of work in quantum cryptography dealing with revoking and certifiably deleting states in the form of ciphertexts and programs [127, 40, 70, 17, 82, 90, 23]. As part of a single unified framework which we call *revocable cryptography*, we show how to use quantum information to revoke large classes of programs, encrypted data, and even cryptographic keys under standard cryptographic assumptions, such as the worst-case hardness of lattice problems.

Outline

Let us now give a brief overview over each chapter and its contributions. We begin with Chapter 2, where we introduce some relevant background on quantum computing and lattices. The latter will especially be relevant for Chapter 4 and Chapter 5.

Revocable programs. In Chapter 3, we focus on the task of revoking programs. Here, we mainly consider the notion of *secure software leasing* (SSL) which was proposed by Ananth and La Placa [17] and captures the following scenario: an authority wishes to "lease" a program f (in the form of a quantum state ϱ_f) to a user who is supposed to "return" the program at a later point in time. Once the supposed copy is returned and verified, the security property requires that the recipient can no longer compute f. Our contributions in this chapter are the following. First, we introduce a new operational security definition for SSL (Section 3.5) by means of a cryptographic security game which does not limit the adversary to performing the honest evaluation procedure. This allows us to significantly strengthen the original security definition introduced by Ananth and La Placa [17]. Second, we give an affirmative answer to a question which was posed by the authors; namely, is it possible to construct an SSL scheme for a simple class of programs from standard cryptographic assumptions? Our main result is an SSL scheme for a large class of evasive functions known as compute-and-compare programs. Here, we consider programs $\mathsf{CC}[f, y]$ which are specified by a function f and a string y within its range: on input x, $\mathsf{CC}[f, y]$ outputs 1, if $f(x) = y$, and 0 otherwise. Our construction is based on Wiesner's conjugate coding technique, and can be instantiated with any cryptographic hash function. To prove the security of our SSL scheme, we have to resolve several technical hurdles; in particular, we have to show that the *monogamy of entanglement* persists, even if the adversary is allowed to interact with a random oracle (which may reveal additional information about the underlying quantum state). Finally, as a complementary result, we make a conceptual connection between unclonable encryption and quantum copy-protection; specifically, we show that we can generically convert any unclonable encryption scheme into a quantum copy-protection scheme for multi-bit point functions, provided it has a mechanism for *wrong-key detection*. We observe that the latter property can easily be achieved by simply outputting a hash of the secret key.

Encryption with publicly-verifiable deletion. In Chapter 4, we focus on the problem of revoking encrypted data. Our results build on a cryptographic notion called *certified deletion*, which was proposed by Broadbent and Islam [40]. While this remarkable notion allows one to certify that a quantum ciphertext was deleted by an untrusted party, it offers no additional functionality. The following question, in particular, was left as open problem: can we enable a cloud server to compute on encrypted data, while also allowing the server to prove data deletion to a client? It is not obvious

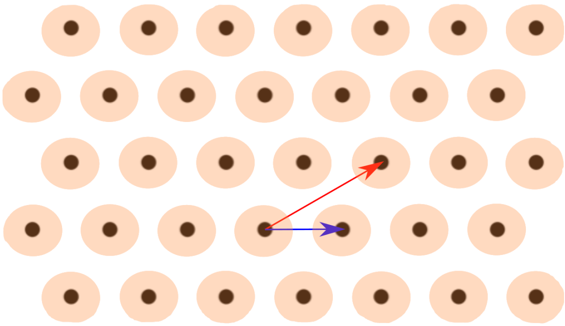

Figure 1.1: Primal Gaussian state.

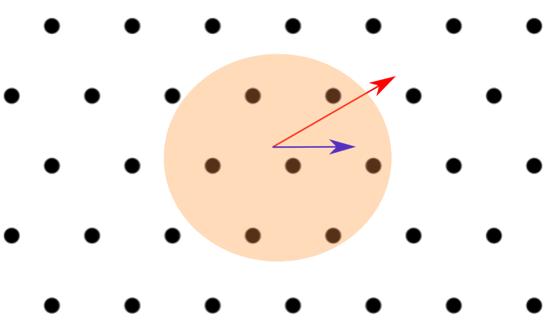

Figure 1.2: Dual Gaussian state.

that such a primitive even exists; for all we know, the server could just homomorphically compute a (secret) classical copy of the encrypted data. We give an affirmative answer to this question and construct the first fully homomorphic encryption scheme with certified deletion, assuming the subexponential hardness of learning with errors [112]. Our scheme has the desirable property that verification of a deletion certificate is completely *public*; meaning that anyone can verify whether deletion has taken place. Central to our construction is the (classical) Dual-Regev encryption scheme (and its variants), which was introduced by Gentry, Peikert and Vaikuntanathan [68]. Our techniques for constructing encryption schemes with certified deletion deviate significantly from the conjugate coding approach used by Broadbent and Islam [40]. Inspired by Regev's reduction from worst-case lattice problems [112], we make use of so-called Gaussian superpositions over lattices, and apply them in the context of certified deletion. Our ciphertext consists of a superposition of Gaussian balls around "random" lattice points. We refer to this superposition as the *primal state* (see Figure 1.1). Depending on whether we encrypt $b = 0$ or $b = 1$, we additionally shift each noisy lattice point by an appropriate vector of large norm. By the learning with errors assumption, this computationally *hides* b. Moreover, using trapdoor information, it is also possible to detect whether such a shift has occurred—thereby allowing one to "decrypt" and to recover the original bit. To enable certified deletion, we make use of the rich structure of Gaussian superpositions. We observe that, when applying the quantum Fourier transform to the primal state, we obtain a superposition over Gaussian-weighted vectors in the *dual lattice*. We call this superposition the *dual state* (see Figure 1.2). Moreover, by equipping the primal state with an appropriate complex phase, we can additionally guarantee that a measurement of the dual state produces a short vector in a *shift* of the dual lattice. We then ask: can such a short vector serve as a deletion certificate? At first sight, it seems as if the *principle of complementarity* in quantum mechanics would immediately prevent an adversary from being able to measure such a quantum state in two incompatible bases, say the computational basis and the Fourier basis. In our case, however, we are dealing with computational assumptions which further complicates the matter. We introduce several new proof techniques that generalize the notion of *collapsing hashes* [125] and allow us to prove a strong notion of certified deletion. Finally, as a simple extension of our homomorphic encryption scheme, we describe a four-message protocol for FHE with simultaneous data deletion, which allows an untrusted quantum server to compute on encrypted data and to simultaneously prove data deletion to a client—all in a single interactive protocol.

Revoking cryptographic keys. In Chapter 5, we build on the no-cloning principle of quantum mechanics and design cryptosystems with key-revocation capabilities. Our contributions are the following. First, we present formal definitions of what it means to "return" a cryptographic key. Broadly speaking, our security notion guarantees that, once the secret key (in the form of quantum states) is revoked from a user, they no longer have the ability to perform the same functionality as before. Second, we construct several fundamental cryptographic primitives with key-revocation capabilities, namely pseudorandom functions, secret-key and public-key encryption, and even fully homomorphic encryption, assuming the hardness of lattice problems. To this end, we adopt many of the techniques we already used in Chapter 4. In particular, we use Gaussian superpositions to generate *quantum decryption keys* which are naturally compatible with the Dual-Regev public-key encryption scheme. To prove the revocation security of our schemes, we have to overcome multiple technical hurdles. First, how can we efficiently check whether a state corresponds to a particular Gaussian superposition? This task is notoriously difficult and has remained a major bottleneck in previous attempts at constructing public-key quantum money schemes from lattices. Fortunately, in the context of key-revocation, we can perform such a verification check in private using appropriate trapdoor information (which is not known to the recipient of the decryption key). We construct an algorithm that allows one to *project* onto particular Gaussian state with access to a lattice trapdoor. Our procedure can be thought of as an explicit quantum reduction between the inhomogenous short integer solution problem [9] and the learning with errors problem [112]. Second, how can we use an adversary that can simultaneously pass revocation and still retain decryption privileges in order to break a computational assumption? During the reduction, we must necessarily simulate the entire security experiment; this includes the revocation phase as well. However, checking whether the returned state is valid requires a trapdoor which is not available during the reduction. It appears the security proof is stuck. We show how to overcome this barrier using techniques from the theory of *quantum rewinding* [100, 138]. Our main result is a simultaneous search-to-decision reduction with quantum auxiliary input, which is tailored towards the Dual-Regev scheme. Informally, our theorem says the following: any strategy that passes revocation (with overwhelming probability) and simultaneously retains its decryption privileges can be converted into an efficient extractor that can "extract" a decryption key from the adversary's state. This crucial insight allows us to complete the reduction, and to base the security of our schemes on the hardness of the short integer solution and learning with errors problems—provided revocation succeeds with high probability.

Chapter 2

PRELIMINARIES

2.1 Notation.

We write $\mathsf{negl}(\cdot)$ to denote any *negligible* function, which is a non-negative function f with the property that, for every constant $c \in \mathbb{N}$, there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$. Similarly, we write $\mathsf{poly}(\cdot)$ to denote any polynomially bounded function f such that $f(n) < n^c$.

We use the following norms:

- For $\mathbf{x} \in \mathbb{C}^n$, we denote the ℓ^2 norm by $\|\mathbf{x}\|$.
- For $\mathbf{M} \in \mathbb{C}^{n\times m}$, we denote by $\|\mathbf{M}\|$ the ℓ_2 norm of the longest column of $\mathbf{M}$.
- For $\mathbf{M} \in \mathbb{C}^{n\times m}$, we denote by $\|\mathbf{M}\|_2 = \sup_{\|\mathbf{x}\|=1} \|\mathbf{M}\mathbf{x}\|$ the operator norm.
- For $\mathbf{M} \in \mathbb{C}^{n\times m}$, we denote the trace norm by $\|\mathbf{M}\|_1 = \mathrm{Tr}[\sqrt{\mathbf{M}^\dagger \mathbf{M}}]$.

The Hellinger distance between two discrete probability distributions P and Q over a finite domain Ω is defined as the quantity,

$$H^2(P,Q) = 1 - \sum_{\omega\in\Omega} \sqrt{P(\omega)Q(\omega)}.$$

The total variation distance between two random variables X and Y with domain Ω is defined as

$$\|X - Y\|_{\mathsf{TV}} = \frac{1}{2}\sum_{\omega\in\Omega} |\Pr[X=\omega] - \Pr[Y=\omega]|.$$

We denote the expectation value of a random variable X with domain $\mathcal{X}$ by

$$\mathbb{E}[X] = \sum_{x\in\mathcal{X}} x \Pr[X=x].$$

The notation $x \overset{\$}{\leftarrow} \Omega$ denotes sampling of x uniformly at random from a domain Ω, whereas $x \sim D$ denotes sampling of an element x according to the distribution D.

Given $m \in \mathbb{N}$ and an integer modulus $q \geq 2$, we represent elements in the ring $\mathbb{Z}_q^m$ as integers in the range $\mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$. Let $p \in \mathbb{N}$. The rounding operation for $q \geq p \geq 2$ is the function

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

2.2 Quantum Computation

For a comprehensive overview of quantum computation, we refer to the introductory texts [105, 133]. 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 $\mathcal{H} = \mathbb{C}^d$. For brevity, we sometimes write $\mathcal{H}_d^n = \mathcal{H}_d^{\otimes n}$, where $\mathcal{H}_d$ is d-dimensional. We use the word *quantum state* to refer to both pure states (unit vectors $|\psi\rangle \in \mathcal{H}$) and density matrices $\varrho \in \mathcal{D}(\mathcal{H})$, 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 convenience, 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 $\| \varrho - \sigma \|_{\mathrm{tr}} \leq \varepsilon$. 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 $\mathcal{H}_B$. 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 $\Phi \otimes \mathbb{1}_R$ 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. A *unitary* $U : L(\mathcal{H}_A) \to L(\mathcal{H}_A)$ is a special case of a quantum channel that satisfies $U^\dagger U = U U^\dagger = \mathbb{1}_A$. An isometry is a linear map $V : L(\mathcal{H}_A) \to L(\mathcal{H}_B)$ with $\dim(\mathcal{H}_B) \geq \dim(\mathcal{H}_A)$ and $V^\dagger V = \mathbb{1}$. A *projector* $\mathbf{\Pi}$ is a Hermitian operator such that $\mathbf{\Pi}^2 = \mathbf{\Pi}$, and a *projective measurement* is a collection of projectors $\{\mathbf{\Pi}_i\}_i$ such that $\sum_i \mathbf{\Pi}_i = \mathbb{1}$. A positive-operator valued measure (POVM) is a set of

Hermitian positive semidefinite operators $\{\mathbf{M}_i\}$ acting on a Hilbert space $\mathcal{H}$ such that $\sum_i \mathbf{M}_i = \mathbb{1}$. The diamond norm of a quantum channel $\Phi : L(\mathcal{H}_A) \to L(\mathcal{H}_B)$ is defined by

$$\|\Phi\|_\diamond = \max_{\varrho} \|(\Phi_{A\to B} \otimes \mathbb{1}_R)(\varrho)\|_1,$$

where the maximization is over all $\varrho \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_R)$ and where R is an arbitrary register.

Quantum algorithms. By a polynomial-time *quantum algorithm* (or QPT algorithm) we mean a polynomial-time uniform family of quantum circuits given by $C = \bigcup_{n\in\mathbb{N}} C_n$, where each circuit $C \in C$ 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 quantum states, [129]). *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* $\lambda \in \mathbb{N}$,

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

We frequently use the following lemma.

Lemma 1 ("Almost As Good As New" Lemma, [3]). *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 then performing the 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,*

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

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

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

The following lemma is a quantum analogue of the standard union bound.

Lemma 2 (Quantum Union Bound, [62])**.** *Let $\varrho \in \mathcal{D}(\mathcal{H})$ be a state and let $\mathbf{\Pi}_1, \dots, \mathbf{\Pi}_n \geq 0$ be sequence of (orthogonal) projections acting on $\mathcal{H}$. Suppose that, for every $i \in [n]$, it holds that* $\mathrm{Tr}[\mathbf{\Pi}_i \varrho] = 1 - \varepsilon_i$, *for $\varepsilon_i \in [0,1]$. Then, if we sequentially measure ϱ with projective measurements $\{\mathbf{\Pi}_1, \mathbf{I} - \mathbf{\Pi}_1\}, \dots, \{\mathbf{\Pi}_n, \mathbf{I} - \mathbf{\Pi}_n\}$, the probability that all measurements succeed is at least*

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

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

Lemma 3 ([**Unruh2013**], Lemma 10)**.** *Let $\mathbf{\Pi}$ be an arbitrary projector and let ϱ be density matrix with* $\mathrm{Tr}[\mathbf{\Pi}\varrho] = 1 - \varepsilon$, *for some $\varepsilon \geq 0$. Then, there exists an ideal state ϱ^{id} with the properties that*

- $\|\varrho - \varrho^{id}\|_{\mathrm{tr}} \leq \sqrt{\epsilon}$
- *ϱ^{id} is a mixture in the image of $\mathbf{\Pi}$, i.e., $\varrho^{id} = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ is a normalized state with $|\psi_i\rangle \in \mathsf{im}(\mathbf{\Pi})$, $\sum_i p_i = 1$ and $p_i \geq 0$, for all i.*

In other words, ϱ^{id} is within trace distance $\varepsilon \geq 0$ of the state ϱ and lies in the image of $\mathbf{\Pi}$.

2.3 Classical and Quantum Entropies

We introduce a few basic notions of entropy – both in the classical and and the quantum setting.

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, [80])**.** *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 3 (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*

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

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 4 (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*

$$p_{\text{guess}}(X|B)_{\varrho} = \sup_{\mathbf{M}_x} \sum_{x \in \mathcal{X}} \Pr[X = x]_{\varrho} \cdot \operatorname{Tr}\left[\mathbf{M}_x \varrho_B\right],$$

where $\{\mathbf{M}_x\}_{x \in \mathcal{X}}$ is a POVM *acting on $\mathcal{H}_B$.*

The following operational meaning of min-entropy is due to Koenig, Renner and Schaffner [92].

Theorem 1 ([92], 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*

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

2.4 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 $\hat{f} : \mathbb{Z}_q^m \to \mathbb{C}$ (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 $\mathbb{Z}_q^m$ is defined by the operation,

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

It is well known that the q-ary quantum Fourier transform can be efficiently performed on a quantum computer for any modulus $q \geq 2$ [79]. Note the quantum Fourier transform of a normalized quantum state

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

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 $x, y \in \mathbb{Z}_q$:*

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

2.5 Generalized Pauli Operators

Definition 5 (Generalized Pauli operators)**.** *Let $q \geq 2$ be a 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*

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

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 6 (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-$\mathbf{Z}$ 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-$\mathbf{Z}$ channel for which $\mathbf{p}$ is the uniform distribution over $\mathbb{Z}_q^m$.

The following well-known lemma states 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. For completeness, we give a proof of the statement below.

Lemma 7 (Pauli-Z twirl)**.** *Let $m, q \in \mathbb{N}$. Then, the uniform Pauli-$\mathbf{Z}$ 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

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

□

2.6 Lattices and the Gaussian Mass

A *lattice* $\Lambda \subset \mathbb{R}^m$ is a discrete subgroup of $\mathbb{R}^m$. We will exclusively consider integer lattices $\Lambda \subseteq \mathbb{Z}^m$ throughout this thesis. 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.

q-ary lattices. In this thesis, 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

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

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

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

where $\mathbf{u} \in \mathbb{Z}^m$ is an arbitrary integer solution to the equation $\mathbf{A}\mathbf{u} = \mathbf{y} \ (\mathrm{mod}\ q)$.

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

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

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

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

Whenever $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ is full-rank, i.e., the subset-sums of the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$, then $\det(\Lambda_q^{\perp}(\mathbf{A})) = q^n$. We use the following facts due to Gentry, Peikert and Vaikuntanathan [68].

Lemma 8 ([68], 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.*

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

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

A simple calculation shows that the Fourier transform of the Gaussian measure is another Gaussian with $\widehat{\varrho_\sigma}(\mathbf{x}) = \sigma^m \varrho_{1/\sigma}(\mathbf{x})$. The *Gaussian mass* of $\Lambda - \mathbf{t}$ is defined as the quantity

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

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

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

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

The following lemma follows from [106, Lemma 2.11] and [68, Lemma 5.3].

Lemma 9 ([68], 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 $\mathbf{e} \sim D_{\mathbb{Z}^m,\sigma}$.*

Lemma 10. *Let $n \in \mathbb{N}$ and let q be a prime with $m \geq 2n \log q$. Let $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$. Let $\mathbf{y} \in \mathbb{Z}_q^n$ be arbitrary. Then, for any $\sigma \geq \omega(\sqrt{\log m})$, there exists a negligible function $\varepsilon(m)$ such that*

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

We make use of the following tail bound for the Gaussian mass of a lattice [20, Lemma 1.5 (ii)].

Lemma 11. *For any m-dimensional lattice Λ and shift $\mathbf{t} \in \mathbb{R}^m$ and for all $\sigma > 0$, $c \geq (2\pi)^{-\frac{1}{2}}$ it holds that*

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

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

The following lemma is a consequence of [101, Lemma 4.4] and [68, Lemma 5.3].

Lemma 12. *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 $\mathbf{y} \in \mathbb{Z}_q^n$:*

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

Definition 7 (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 complex-valued function $f : \mathbb{Z}^m \to \mathbb{C}$ and integer lattice $\Lambda \subseteq \mathbb{Z}^m$, the well-known *Poisson summation formula* relates $f(\Lambda)$ to its Fourier transform $\hat{f}$ over the dual lattice, i.e.,

$$f(\Lambda) = \det(\Lambda^*)\hat{f}(\Lambda^*).$$

We use the following variant of the formula which applies to Gaussians and q-ary lattices.

Lemma 13 (Poisson summation for Gaussians over q-ary lattices)**.** *Let q 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 $\sigma > 0$. Then,*

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

Proof. Because $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ is a full-rank matrix, it holds that $\det(\Lambda_q^\perp(\mathbf{A})) = q^n$. Let $\Lambda_q^{\mathbf{v}}(\mathbf{A})$ be the lattice coset given by $\Lambda_q^\perp(\mathbf{A}) + \mathbf{u}$, for some arbitrary solution $\mathbf{u} \in \mathbb{Z}^m$ with $\mathbf{A} \cdot \mathbf{u} = \mathbf{v} \pmod{q}$. Recall that the Fourier transform of the Gaussian measure satisfies

$$\widehat{\varrho_\sigma}(\mathbf{x}) = \sigma^m \varrho_{1/\sigma}(\mathbf{x}), \quad \forall \mathbf{x} \in \mathbb{R}^m.$$

Therefore, it follows from the Poisson summation formula that

$$\begin{aligned}\sum_{\mathbf{x}\in\Lambda_q^{\mathbf{v}}(\mathbf{A})} \varrho_\sigma(\mathbf{x})\cdot e^{-\frac{2\pi i}{q}\langle\mathbf{w},\mathbf{x}\rangle} &= \sum_{\mathbf{x}\in\Lambda_q^{\perp}(\mathbf{A})+\mathbf{u}} \varrho_\sigma(\mathbf{x})\cdot e^{-\frac{2\pi i}{q}\langle\mathbf{w},\mathbf{x}\rangle}\\ &= \frac{\sigma^m}{\det(\Lambda_q^{\perp}(\mathbf{A}))} \sum_{\mathbf{y}\in\frac{1}{q}\Lambda_q(\mathbf{A})} \varrho_{1/\sigma}(\mathbf{w}+q\cdot\mathbf{y})\cdot e^{2\pi i\langle\mathbf{y},\mathbf{u}\rangle}\\ &= \frac{\sigma^m}{q^n} \sum_{\mathbf{y}\in\Lambda_q(\mathbf{A})} \varrho_{q/\sigma}(\mathbf{w}+\mathbf{y})\cdot e^{\frac{2\pi i}{q}\langle\mathbf{y},\mathbf{u}\rangle}\\ &= \frac{\sigma^m}{q^n} \sum_{\mathbf{y}\in\mathbb{Z}_q^n} \varrho_{q/\sigma}(\mathbf{w}+\mathbf{A}^\intercal\mathbf{y}+q\cdot\mathbb{Z}^m)\cdot e^{\frac{2\pi i}{q}\langle\mathbf{A}^\intercal\mathbf{y},\mathbf{u}\rangle}\\ &= \frac{\sigma^m}{q^n}\cdot \sum_{\mathbf{y}\in\mathbb{Z}_q^n} \varrho_{q/\sigma,q}(\mathbf{w}+\mathbf{A}^\intercal\mathbf{y})\cdot e^{\frac{2\pi i}{q}\langle\mathbf{y},\mathbf{v}\rangle}.\end{aligned}$$

□

We use the following lemma due to Brakerski [36] which 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 14 ([36])**.** *Let $q \geq 2$ be a modulus and $\mathbf{x} \in \mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$. Let $\sigma > 0$. Then,*

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

A consequence of the tail bound in Lemma 11 is that the Gaussian distribution $D_{\mathbb{Z}^m,\sigma}$ is essentially only supported on the finite set $\{\mathbf{x} \in \mathbb{Z}^m : \|\mathbf{x}\| \leq \sigma\sqrt{m}\}$, which suggests the use of *truncation*.

Definition 8 (Truncated discrete Gaussian distribution)**.** *Let $m \in \mathbb{N}$, $q \geq 2$ be an integer modulus and let $\sigma > 0$ be a parameter. Then, the* truncated *discrete Gaussian distribution $D_{\mathbb{Z}_q^m,\sigma}$ with finite support $\{\mathbf{x} \in \mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m : \|\mathbf{x}\| \leq \sigma\sqrt{m}\}$ is defined as the density*

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

We use the following *noise smudging* property of the discrete Gaussian.

Lemma 15 (Noise smudging, [57])**.** *Let $y, \sigma > 0$. Then, the total variation distance between the distribution $D_{\mathbb{Z},\sigma}$ and $D_{\mathbb{Z},\sigma+y}$ is at most y/σ.*

Occasionally, we also use the following variant of noise smudging which 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 16 ([38], Lemma 2.4)**.** *Let $q \geq 2$ be a modulus, $m \in \mathbb{N}$ and $\sigma > 0$. Then, for any $\mathbf{e}_0 \in \mathbb{Z}^m$,*

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

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

Lemma 17. *Let $n \in \mathbb{N}$ and let q be a prime with $m \geq 2n\log q$. Let $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$. Then, for any $\sigma \geq \omega(\sqrt{\log m})$, there exists a negligible $\varepsilon(m)$ such that*

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

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

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

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

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

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

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

we get that

$$\begin{aligned} &\max_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\substack{\mathbf{x} \in \mathbb{Z}_q^m,\, \|\mathbf{x}\| \le \sigma\sqrt{m} \\ \mathbf{A}\mathbf{x} = \mathbf{y} \pmod{q}}} \left\{ \frac{\varrho_\sigma(\mathbf{x})}{\displaystyle\sum_{\substack{\mathbf{z} \in \mathbb{Z}_q^m, \|\mathbf{z}\| \le \sigma\sqrt{m} \\ \mathbf{A}\mathbf{z} = \mathbf{y} \pmod{q}}} \varrho_\sigma(\mathbf{z})} \right\} \\ &\le \max_{\mathbf{y} \in \mathbb{Z}_q^n} \left\{ \sup_{\mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A})} D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}(\mathbf{x}) + \Pr_{\mathbf{v} \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\sigma}} \left[\|\mathbf{v}\| > \sigma\sqrt{m} \right] \right\}. \end{aligned}$$

Because $\sigma \ge \omega(\sqrt{\log m})$, the claim then follows from Lemma 10 and Lemma 12. □

2.7 Cryptography

In this section, we review several definitions in cryptography.

Public-key encryption

Definition 9 (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* QPT *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* 1^λ *and outputs a public key* pk *and secret key* sk.

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

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

Definition 10 (Correctness of PKE)**.** *For any* $\lambda \in \mathbb{N}$*, and for any* $m \in \mathcal{M}$*:*

$$\Pr\left[\mathsf{Dec}(\mathsf{sk}, \mathsf{CT}) \ne m \;\middle|\; \substack{(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{pk}, m)} \right] \le \mathsf{negl}(\lambda).$$

Definition 11 (IND-CPA security)**.** *Let* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ *be a* PKE *scheme and* $\mathcal{A}$ *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 $\mathsf{IND\text{-}CPA}$-secure if, for any QPT adversary $\mathcal{A}$, it holds that

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

The Short Integer Solution problem

The (inhomogenous) SIS problem was introduced by Ajtai [9] in his seminal work on average-case lattice problems. The problem is defined as follows.

Definition 12 (Inhomogenous SIS problem,[9]). *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 $\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod q$ 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 [103] showed that the SIS problem is, on the average, as hard as approximating worst-case lattice problems to within small factors. Subsequently, Gentry, Peikert and Vaikuntanathan [68] gave an improved reduction showing that, for parameters $m = \mathsf{poly}(n)$, $\beta = \mathsf{poly}(n)$ and prime $q \geq \beta \cdot \omega(\sqrt{n \log q})$, the average-case $\mathsf{SIS}^m_{n,q,\beta}$ problem is as hard as approximating the shortest independent vector problem (SIVP) problem in the worst case to within a factor $\gamma = \beta \cdot \tilde{O}(\sqrt{n})$. We assume that $\mathsf{SIS}^m_{n,q,\beta}$, for $m = \Omega(n \log q)$, $\beta = 2^{o(n)}$ and $q = 2^{o(n)}$, is hard against quantum adversaries running in time $\mathsf{poly}(q)$ with success probability $\mathsf{poly}(1/q)$.

The Learning with Errors problem

The *Learning with Errors* problem was introduced by Regev [112] and serves as the primary basis of hardness of post-quantum cryptosystems. The problem is defined as follows.

Definition 13 ("Search" LWE, [112]). *Let $n, m \in \mathbb{N}$, 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 $(\mathbf{A}, \mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q)$ 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 14 ("Decisional" LWE, [112]). *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*

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

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

As shown in [112], the $\mathsf{LWE}_{n,q,\alpha q}^m$ 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 $\mathsf{LWE}_{n,q,\alpha q}^m$ which relies on the worst case hardness of approximating short vector problems in lattices to within a subexponential factor. We assume that the $\mathsf{LWE}_{n,q,\alpha q}^m$ problem, for $m = \Omega(n \log q)$, $q = 2^{o(n)}$, $\alpha = 1/2^{o(n)}$, is hard against quantum adversaries running in time $\mathsf{poly}(q)$. We note that this parameter regime implies $\mathsf{SIS}_{n,q,\beta}^m$ [120].

Trapdoors for lattices

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

Theorem 2 ([102], Theorem 5.1)**.** *Let* $n, m \in \mathbb{N}$ *and* $q \in \mathbb{N}$ *be a prime with* $m = \Omega(n \log q)$*. There exists a randomized algorithms with the following properties:*

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

Chapter 3

REVOKING UNCLONABLE PROGRAMS

In this chapter, we investigate the following question: can we use the no-cloning principle of quantum mechanics and encode a program in such a way that it can be evaluated, yet it cannot be *pirated*? Naturally, we would also like to ensure that, once the program is "returned," the recipient loses its ability to evaluate it. Our main result is secure software leasing (SSL) scheme for a large class of evasive functions known as compute-and-compare programs.

Organization. First, we focus our attention on the notion of quantum copy-protection which was proposed by Aaronson [2] as a means of software protection. In Section 3.3, we present a formal definition of what a copy-protection scheme is. Then, in Section 3.4, we construct a quantum copy-protection scheme for multi-bit point functions. In the second half of the chapter, we consider the weaker notion of secure software leasing (SSL) which was proposed by Ananth and La Placa [17]. In Section 3.5, we give a formal defintion of what an SSL scheme is. Then, in Section 3.6, we construct an SSL scheme for single-bit point functions. Finally, in Section 3.6, we prove our main result; namely, we give an SSL scheme for compute-and-compare programs as a simple extension of our SSL scheme for single-bit point functions.

3.1 Introduction

Aaronson [2] initiated the formal study of quantum copy-protection schemes, and speculated that quantum cryptography could offer a solution to software piracy thanks to the no-cloning theorem. Copy-protection captures the following cryptographic task. A vendor wishes to encode a program in such a way that a user who receives the encoded program is able to run it on arbitrary inputs. However, the recipient should not be able to create functionally equivalent "pirated" copies of the original program. More concretely, no user should be able to process the encoded program so as to split it into two parts, each of which allows for the evaluation of the function implemented by the original program. Rigorous copy-protection of any kind is trivially impossible to achieve classically. This is because any information that the user receives can simply be copied. In the quantum realm, however, the no-cloning theorem prevents any naive copying strategy from working unconditionally, and copy-protection seems, at least in principle, possible. The key question then becomes: *Is it possible to encode functionality into a quantum state while at the same time preserving the no-cloning property?*

To be precise, we are not satisfied with preventing an adversary from copying the state that encodes the program (this is certainly a necessary condition), but we also require that there is no other way for a (computationally bounded) user to process the state into two parts (not necessarily a copy of the original) so as to allow each half to recover the input-output behaviour of the encoded program.

Quantum copy-protection was first formalized by Aaronson [2]. One of the first observations there is that families of *learnable* functions cannot be copy-protected: access to a copy-protected program, and hence its input-output behaviour, allows one to recover a classical description of the program itself, which can be copied. In [2], Aaronson provides some formal definitions and constructions of copy-protection schemes. More precisely, Aaronson describes:

- A provably secure scheme to copy-protect any family of efficiently computable functions which is not quantumly learnable, assuming a *quantum* oracle implementing a certain family of unitaries.
- Two candidate schemes to copy-protect point functions in the plain model, although neither of the two features a proof of security.

In recent work [5], Aaronson et al provide a scheme to copy-protect any family of efficiently computable functions which is not quantumly learnable, assuming access to a *classical* oracle, i.e., an oracle (which can be queried in superposition) that implements a classical function. We emphasize, however, that this classical function is *dependent* on the function that one wishes to copy-protect. In particular, the oracle is impossible to realize in general, as it implies an ideal obfuscator for the function f that is being copy-protected, and is thus very strong. In particular, the following questions were left open in [2]: Does there exist a scheme to copy-protect any non-trivial family of functions (the simplest example being point functions) with provable security in the plain model using standard assumptions? What about larger classes of programs?

On the negative side, aside from the impossibility of copy protecting families of learnable functions, it has remained an open question to determine whether a more general impossibility result applies. In a recent result, Ananth and La Placa [17] prove that a universal copy-protection scheme cannot exist, assuming the quantum hardness of the learning with errors problem [111] and the existence of quantum fully homomorphic encryption.

On top of proving the impossibility of a general copy-protection scheme for all unlearnable functions, Ananth and La Placa introduce in [17] a weaker notion of copy-protection, which they call "secure software leasing" (SSL). The sense in which the latter is weaker than copy-protection is that one assumes that the freeloaders $\mathcal{B}$ and C (now a single adversary) are limited to performing the honest evaluation procedure only. Rather than emphasizing the impossibility of simultaneous evaluation on inputs chosen by a challenger, SSL captures the essence of quantum copy-protection in the following scenario. An authority (the lessor) wishes to lease a copy ϱ_f of a function $f \in \mathcal{F}_\lambda$ to a user (the lessee) who is supposed to return back ϱ_f at a later point in time, as specified by the lease agreement. Once the program copy is "revoked" and verified by the lessor, the security property requires that the adversary can no longer compute f. More formally, no adversary should be able to produce a (possibly entangled) quantum state such that:

- One half of the state is deemed valid by the lessor, once it is returned.
- The other half can be used to honestly evaluate f on every input of the adversary's choosing.

Surprisingly, Ananth and La Placa were able to show in [17] that a general SSL scheme is also impossible, despite having weaker security requirements compared to copy-protection. On the positive side, the authors describe an SSL scheme for general evasive circuits assuming the existence of subspace-hiding obfuscators [137] and the quantum hardness of the learning with errors problem [111]. Because subspace-hiding obfuscators are only known to exist under indistinguishability obfuscation [137, 117], the same applies to the security of the scheme proposed in [17]. A key question, in particular, which their work left open is the possibility that one can construct SSL for more primitive classes of programs under standard cryptographic assumptions.

Our contributions

Let us now give an overview of our results.

Quantum copy-protection. We approach the task of quantum copy-protection from the positive side; specifically, we give a copy-protection scheme for multi-bit point functions and we prove its security in the so-called quantum random oracle model – a standard cryptographic assumption. Our construction can be instantiated with any cryptographic hash function, for example using SHA-3.

A desirable feature of our scheme is that the copy-protected program does not involve multi-qubit entanglement – in fact it only involves BB84 states and computational and Hadamard basis measurements. This is in contrast to previous candidate schemes for point functions in [2], whose security is only conjectured, and which employ highly entangled states. The simple structure of the copy-protected program is advantageous for, e.g., error-corrected storage of the copy-protected program. We point out, however, that in a practical implementation of our scheme, where the oracle is replaced by a hash function, evaluation of the copy-protected program on an input requires *coherently* computing the hash function in an auxiliary register. This operation requires universal quantum computation.

Our scheme is not in the plain model. The (quantum) random oracle model, however, enjoys widespread acceptance and popularity in (post-quantum/quantum) cryptography, and many schemes designed for, and deployed in, practical applications enjoy provable security in that model only. Our security definition is essentially analogous to the original definition in [2] but differs more significantly from the more recent definition in [5], which is weaker. In Section 3.3, we give a more detailed comparison of our definition with the ones in [2] and [5].

Our techniques and construction are inspired by recent work on *unclonable encryption* by Broadbent and Lord [41]. The main technical ingredient on which their construction relies are *monogamy of entanglement games*, introduced and studied extensively in [123], which they combine with an adaption of the one-way-to-hiding (O2H) lemma of [127] for a security analysis in the quantum random oracle model. In a nutshell, (a special case of) the latter lemma allows one to upper bound the probability that an algorithm outputs $H(x)$, where H is a random oracle and x is any string in the domain, in terms of the probability that the algorithm "queries" at x at some point during its execution. The adaption of [41] extends the applicability of the O2H lemma to a setting that involves *two players*, and upper bounds the probability that the two (possibly entangled) players *simultaneously* guess $H(x)$ by the probability that they both query at x at some point during the execution of their respective strategies.

A sketch of our copy-protection scheme

Our quantum copy-protection scheme allows a software vendor to encode a multi-bit point function in such a way that it can be evaluated on any input, yet it cannot be pirated. More specifically, we consider the class of functions of the form $P_{y,m}$, for some strings $y, m \in \{0,1\}^\lambda$ with

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

Multi-bit point functions can potentially serve as *password authentication* programs, since the recipient of the program can easily check whether a given input matches a hidden password y, and additionally learn a message m if the password is correct. Naturally, we require that both the password and the message have sufficient amounts of entropy.

Our construction is inspired by recent work on *unclonable encryption* by Broadbent and Lord [41] which revisits the cryptographic notion first proposed by Gottesman [75]. In an unclonable encryption scheme, one encrypts a *classical* message in a *quantum* ciphertext, in such a way that the latter cannot be processed and *split* into two parts such that each half, together with a classical secret key, enables decryption. The setting of unclonable encryption is very similar to that of copy-protection, the main difference being that there is no "functionality" associated to the quantum ciphertext, other than it being used for recovering the encrypted message.

Our simple observation is that any *unclonable encryption* scheme can be generically turned into a quantum copy-protection scheme for multi-bit point functions as follows. To copy-protect a function $P_{y,m}$, simply encrypt the message m with the secret key y. Then, provided there exists a mechanism for *wrong-key detection*, this already achieves our goal: to evaluate at point x, attempt to decrypt using x; if decryption succeeds output the decrypted message, if decryption fails, output 0^λ. We observe that any unclonable encryption scheme can be easily upgraded to achieve wrong-key detection in the QROM, thereby yielding the desired copy-protection scheme. Our notion of wrong-key detection for quantum encryption schemes is inspired by the work of Canetti et al. [46] who previously introduced a similar property for classical encryption schemes.

To illustrate how we can construct a copy-protection scheme for multi-bit point functions from any unclonable encryption, we now consider a concrete example. This follows the unclonable encryption scheme from Broadbent and Lord [41] which itself is rooted in Wiesner's conjugate coding scheme [132]. The basic idea is that it is possible to encrypt a message $m \in \{0,1\}^\lambda$ by sampling a random string $r \overset{\$}{\leftarrow} \{0,1\}^\lambda$ and by making λ uniformly random choices of basis (either computational or Hadamard) which we denote by $\theta \in \{0,1\}^\lambda$. Then, one can encode each bit of the one-time padded message $m \oplus r$ either in the computational or the Hadamard basis, according to θ. Formally, letting $k = (r, \theta)$ denote the secret key, this amounts to preparing the following quantum ciphertext on λ many qubits:

$$\mathsf{Enc}_k(m) = |(m \oplus r)^\theta\rangle\langle(m \oplus r)^\theta|\,,$$

where we use the notation $|x^\theta\rangle = H^{\theta_1}|x_1\rangle \otimes \dots H^{\theta_\lambda}|x_\lambda\rangle$ and $|b^s\rangle = H^s|b\rangle$, for $b, s \in \{0,1\}$. Given the key $k = (r, \theta)$ and ciphertext, one can easily "decrypt" and recover the string m by first measuring each qubit of $\mathsf{Enc}_k(m)$ in the basis specified by θ, and then uncomputing the one-time pad specified by r. We now show how to bootstrap such an encryption scheme into a copy-protection scheme for multi-bit point functions as follows. The basic idea is the following. To copy-protect $P_{y,m}$, simply hand out $\mathsf{Enc}_y(m)$ with $y = k = (r, \theta)$ together with some classical information that enables an evaluator to "recognize" an incorrect key. One can take the latter information to be $H(y)$, for some hash function H (or a uniformly random function H, if one works in the random oracle model). Then, to evaluate the program on some input x, the evaluator first checks whether $H(x)$ matches the hash $H(y)$. If not, the evaluator will conclude that the output is 0^λ. Otherwise, if true, the evaluator can simply "decrypt" $\mathsf{Enc}_y(m)$ as before. If the output length of the hash function H is sufficiently large, say 2λ, the resulting scheme achieves the aforementioned wrong-key detection property with overwhelming probability by a standard birthday bound.

Copy-protection security. Before we expand on the technical hurdles we encounter when proving the security of our copy-protection scheme, let us first formalize the property in a bit more detail. We say that a quantum copy-protection scheme is *secure* for a family of functions $\mathcal{F}_\lambda$ (as well as a distribution $\mathcal{D}$ over $\mathcal{F}_\lambda$) if no adversary $(\mathcal{A}, \mathcal{B}, C)$—consisting of a triple of quantum polynomial time algorithms: a "pirate" $\mathcal{A}$, say Alice, and two "freeloaders" $\mathcal{B}$ and C, say Bob and Charlie)—can succeed with sufficiently high (i.e., non-trivial) probability at the following game:

- $\mathcal{A}$ receives a copy-protected program ϱ_f from the challenger (where the program $f \in \mathcal{F}_\lambda$) is sampled from some distribution $\mathcal{D}_{\mathcal{F}_\lambda}$). Next, $\mathcal{A}$ creates a bipartite state on registers B and C, and sends B to $\mathcal{B}$ and C to C.
- The challenger samples a pair (x_1, x_2) of inputs to f from a suitable challenge distribution (which is allowed to depend on f), and sends x_1 to $\mathcal{B}$ and x_2 to C.
- $\mathcal{B}$ and C, who are not allowed to communicate, return bits b_1 and b_2, respectively.
- $(\mathcal{A}, \mathcal{B}, C)$ win if $b_1 = f(x_1)$ and $b_2 = f(x_2)$.

The security of the aforementioned unclonable encryption scheme (and, by implication, the security of our copy-protection scheme) crucially leverages the following property: it is impossible for any pirate who has the ciphertext but does not know θ nor r, to produce a state on two registers BC such that two "freeloaders," say Bob and Charlie, with access to registers B and C, respectively, *as well as* access to θ and r, can simultaneously recover m. Note that the latter property crucially holds even when both Bob and Charlie are *simultaneously* receiving θ and r. This property is essentially a consequence of the *monogamy of entanglement* and is captured formally in [123, 41] via the study of monogamy of entanglement games. In particular, a rephrasing of the results of [123] is that, for any (unbounded) strategy of Alice, Bob and Charlie, the probability that both Bob and Charlie are able to simultaneously recover m is exponentially small in λ.

Unfortunately, our proof of security does not immediately follow from the security of the underlying unclonable encryption scheme, mainly due to the fact that the encoded program $\mathsf{Enc}_y(m)$ also consists of the classical hash $H(y)$ which further complicates the matter. To show security of our scheme, it suffices to argue that the security of the underlying unclonable encryption scheme is preserved, even if the adversary additionally receives as input a classical hash $H(y)$ which depends on the key of the quantum ciphertext. To carry out the security reduction, we use a variant of the so-called one-way-to-hiding (O2H) lemma [127] due to Unruh. This allows us to obtain an upper bound on the probability that an adversary distinguishes $H(y)$ from a uniformly random string in the co-domain of H, in terms of the probability that such an adversary queries the oracle at y.

We prove the following result, which is a consequence of Lemma 18 and Theorem 3.

Theorem (Informal)**.** *Assuming unclonable-secure quantum encryption schemes exist, there exists a quantum copy-protection scheme for multi-point functions which is secure against arbitrary challenge distributions in the quantum random oracle model.*

Next, we consider the notion of secure software leasing.

Secure software leasing

As we mentioned earlier, the original definition of secure software leasing in [17] is a weaker version of copy-protection in the following two ways:

- The lessor performs a prescribed verification procedure on a register returned by the lessee.
- The lessee is required to perform the honest evaluation procedure with respect to any post-verification registers in the lessee's possession.

We revisit the notion of secure software leasing from a similar perspective as in our copy-protection definition. Our main contributions are the following. First, we introduce a new and intuitive SSL definition (Section 3.5) by means of a cryptographic security game which does not limit the adversary to performing the honest evaluation on any post-verification registers.[1] Informally, any SSL scheme $(\mathsf{SSL.Gen}, \mathsf{SSL.Lease}, \mathsf{SSL.Eval}, \mathsf{SSL.Verify})$ according to our definition should satisfy the following property. After receiving a leased copy of f, denoted by ϱ_f (and generated using $\mathsf{SSL.Lease}$), and a circuit for $\mathsf{SSL.Eval}$, no adversary should be able to produce a (possibly entangled) quantum state σ on two registers R_1 and R_2 such that:

- $\mathsf{SSL.Verify}$ deems the contents of register R_1 of $\sigma_{\mathsf{R}_1\mathsf{R}_2}$ to be valid, and
- the adversary can predict the output of circuit f (on challenge inputs chosen by the lessor) using an arbitrary measurement of the post-verification state in register R_2.

Our definition remains faithful to the idea of secure software leasing from [17], while at the same time offering a stronger security guarantee.

Second, we show that our definition of security is achievable with a standard negligible security bound in the quantum random oracle model for the class of compute-and-compare programs [131,

[1] The SSL definition in [17] is not "operational" and cannot be directly phrased as a security game.

76]. A compute-and-compare program $\mathsf{CC}[f, y]$ is specified by an efficiently computable function $f : \{0,1\}^n \to \{0,1\}^m$ and a string $y \in \{0,1\}^m$ in its range, where

$$\mathsf{CC}[f, y](x) = \begin{cases} 1 & \text{if } f(x) = y\,, \\ 0 & \text{if } f(x) \neq y\,. \end{cases}$$

Note that point functions are a special case of compute-and-compare programs where the function f is the identity map. we show how to lease $\mathsf{CC}[f, y]$ in the following simple way: the encoded program consists of (a description of) a function f in the clear, together with a quantumly encoded version of the point function with marked input y. The intuition is that it is enough to protect the marked input y in order to render $\mathsf{CC}[f, y]$ unclonable. At first, it might seem surprising that one can give f in the clear while preserving unclonability, as the encoded program now leaks significantly more information than its input/output behavior alone. At a second thought, however, it is in fact quite natural that one can render a functionality "unclonable" by just making some sufficiently important component of it unclonable. Indeed, it is straightforward to show that the SSL security of the extended construction reduces to the SSL security of the underlying point function scheme.

In Section 3.5 we show the following key property about our SSL scheme in Construction 4: once a leased copy is successfully returned to the lessor, no adversary can distinguish the marked input of a compute-and-compare program from a random (non-marked) input with probability better than $1/2$, except for a negligible advantage. We prove the following in Theorem 5.

Theorem (Informal)**.** *There exists an* SSL *scheme for compute-and-compare programs which is secure against a natural class of input challenge distributions in the quantum random oracle model.*

The result follows from a standard application of the O2H lemma and a particular "uncertainty relation" variant of the monogamy of entanglement property which appeared in a work of Unruh [127]. The latter appears in similar contexts in the quantum key-distribution literature. Note that the technical complications arising in the proof of security of our original copy-protection scheme do not appear in the SSL security proof. Crucially, this is because we can leverage the fact that the lessor is performing a prescribed verification procedure.

Related work

Unclonable encryption. This cryptographic functionality was formalized recently by Broadbent and Lord [41], and informally introduced earlier by Gottesman [75]. In an unclonable encryption scheme, one encrypts a *classical* message in a *quantum* ciphertext, in such a way that the latter cannot be processed and *split* into two parts, each of which, together with a classical secret key, enables decryption. The setting of unclonable encryption is very similar to that of copy-protection, the main difference being that there is no "functionality" associated to the quantum ciphertext, other than it being used for recovering the encrypted message. As we mentioned earlier, our copy-protection scheme is inspired by the unclonable encryption scheme in [41], and our analysis extends some of the techniques developed there.

Revocable quantum timed-release encryption. Timed-release encryption (also known as time-lock puzzles) is an encryption scheme that allows a recipient to decrypt only after a specified amount time, say T, has passed. Unruh [127] gave the first quantum timed-release encryption scheme that is "revocable" in the sense that a user can return the timed-release encryption before time T, thereby losing all access to the data. It is easy to see that this notion is impossible to achieve classically for precisely the same reason copy-protection is impossible: any adversary can simply generate copies of the classical ciphertext or source code, respectively. From a technical point of view, revocable quantum timed-release encryption shares many similarities with the notion of "secure software leasing" in [17]. Besides the fact that the former encodes a *plaintext* and the latter encodes a *program*, the security property essentially remains the same: once a quantum state is returned and successfully verified, the user is supposed to lose all relevant information. Our proof of security for the SSL scheme in Construction 3 is inspired by Unruh's proof for revocable one-way timed-release quantum encryption in [127].

Open questions

Our work is the first to construct a copy-protection scheme in a standard cryptographic model (the QROM). It leaves several questions open. The most pressing ones are the following.

- First, is it possible to extend the security of quantum copy-protection schemes towards multiple copies? In other words, the pirate receives k copy-protected copies of a program, and we ask $k + 1$ freeloaders to succeed. We believe that our scheme can achieve such a notion, but with a security that becomes worse as k grows. Providing a scheme where security does not depend on k is an interesting open question.
- Second, is it possible to remove the requirement of a random oracle, and to achieve a scheme with non-trivial security against malicious adversary in the plain model? We think that this would require fundamentally different techniques.

Subsequent work

We remark that a series of subsequent works have meanwhile improved on some of our results on copy-protection and secure software leasing. Broadbent et al. [42] showed how to construct an *information-theoretic* SSL scheme for compute-and-compare programs from any quantum message authentication code. This improves on our SSL scheme in Construction 3 which relies on the quantum random oracle heuristic. Using *subspace coset states* rather than BB84 states, Ananth et al. [18] obtained a quantum copy-protection for single-bit point function which achieves the standard notion of *negligible adversarial security* in the QROM. This is significantly more challenging to achieve and involves advanced techniques.

3.2 Preliminaries

We now review some technical background which is required for our constructions.

Monogamy of entanglement games

For a detailed introduction to monogamy-of-entanglement games, we refer the reader to the seminal paper on the topic [123], where they were introduced and studied extensively. In this section, we limit ourselves to introducing a version of a monogamy-of-entanglement game that suffices for our purpose. Let $\lambda \in \mathbb{N}$. The game is between a challenger and an adversary, specified by a triple of interactive quantum machines $(\mathcal{A}, \mathcal{B}, C)$ (for a formal definition of interactive quantum machine we refer the reader to [126]). For brevity, we use the notation $|x^\theta\rangle = H^\theta |x\rangle$, where $H^\theta = H^{\theta_1} \otimes \ldots \otimes H^{\theta_\lambda}$ and $\theta, x \in \{0,1\}^\lambda$. The game takes place as follows:

1. The challenger samples $x, \theta \leftarrow \{0,1\}^\lambda$ and sends the state $|x^\theta\rangle$ to $\mathcal{A}$.
2. $\mathcal{A}$ sends a quantum register to $\mathcal{B}$ and one to C.
3. The challenger sends θ to both $\mathcal{B}$ and C.
4. $\mathcal{B}$ and C return strings x' and x'' to the challenger.

The players $\mathcal{A}$, $\mathcal{B}$ and C are not allowed to communicate other than where specified by the game. Finally, $\mathcal{A}, \mathcal{B}, C$ win if $x = x' = x''$.

The following lemma, from [123], upper bounds the winning probability of an adversary in the game. As stated in the form below, this lemma appears in [41].

Lemma 1 ([41], Theorem 1)**.** *Let $\lambda \in \mathbb{N}$ be a parameter. For any Hilbert spaces $\mathcal{H}_\mathsf{B}$ and $\mathcal{H}_\mathsf{C}$, any families of* POVMs *on these Hilbert spaces, respectively,*

$$\left\{\left\{B_\theta^x\right\}_{x\in\{0,1\}^\lambda}\right\}_{\theta\in\{0,1\}^\lambda} \quad \textit{and} \quad \left\{\left\{C_\theta^x\right\}_{x\in\{0,1\}^\lambda}\right\}_{\theta\in\{0,1\}^\lambda},$$

and any CPTP *map $\Phi : \mathcal{D}\left((\mathbb{C}^2)^{\otimes\lambda}\right) \to \mathcal{D}(\mathcal{H}_\mathsf{B} \otimes \mathcal{H}_\mathsf{C})$, we have:*

$$\mathbb{E}_{\theta\in\{0,1\}^\lambda}\mathbb{E}_{x\in\{0,1\}^\lambda}\mathrm{Tr}\Big[(B_\theta^x \otimes C_\theta^x)\,\Phi\left(|x^\theta\rangle\langle x^\theta|\right)\Big] \leq \left(\frac{1}{2} + \frac{1}{2\sqrt{2}}\right)^\lambda. \tag{3.1}$$

The quantum random oracle model

Oracles with quantum access have been studied extensively, for example in [30, 34]. We say that a quantum algorithm $\mathcal{A}$ has oracle access to a classical function $H : \{0,1\}^\lambda \to \{0,1\}^m$, denoted by $\mathcal{A}^H$, if $\mathcal{A}$ is allowed to use a unitary gate O^H at unit cost in time. The unitary O^H acts as follows on the computational basis states of a Hilbert space $\mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B}$ of $\lambda + m$ qubits:

$$O^H : \quad |x\rangle_\mathsf{A} \otimes |y\rangle_\mathsf{B} \longrightarrow |x\rangle_\mathsf{A} \otimes |y \oplus H(x)\rangle_\mathsf{B} ,$$

where the operation $\oplus$ denotes bit-wise addition modulo 2. In general, we can model the interaction of a quantum algorithm that makes q queries to an oracle H as $(UO^H)^q$, i.e., alternating unitary computations and queries to the oracle H, where U is some operator acting on $\mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B} \otimes \mathcal{H}_\mathsf{C}$, where $\mathcal{H}_\mathsf{C}$ is some auxiliary Hilbert space [30, 34, 127][2]. We call a (possibly super-polynomial-time) quantum algorithm $\mathcal{A}$ with access to an oracle O *query-bounded* if $\mathcal{A}$ makes at most polynomially many (in the size of its input) queries to O. The *random oracle* model refers to a setting in which the function $H : \{0,1\}^\lambda \to \{0,1\}^m$ is sampled uniformly at random. Random oracles play an important role in cryptography as models for cryptographic hash functions in the so-called random oracle model (ROM) [26]. For post-quantum and quantum cryptography, random oracles modelling hash functions need to be *quantum* accessible (i.e., accessible as a unitary gate, and thus in superposition), resulting in what is known as the quantum random oracle model (QROM) [34]. Despite being uninstantiable in principle [44, 60], modeling hash functions in the (Q)ROM is considered a standard assumption in cryptography.

Some technical lemmas

Below, we denote by $\mathrm{Bool}(\lambda, m)$ the set of functions from $\{0,1\}^\lambda$ to $\{0,1\}^m$.

Lemma 2. *Let* $f : \mathrm{Bool}(\lambda, m) \to \mathbb{R}$, *and* $x \in \{0,1\}^\lambda$. *For* $H \in \mathrm{Bool}(\lambda, m)$ *and* $y \in \{0,1\}^m$, *let* $H_{x,y} \in \mathrm{Bool}(\lambda, m)$ *be such that*

$$H_{x,y}(s) = \begin{cases} H(s) & \textit{if } s \neq x\,, \\ y & \textit{if } s = x\,. \end{cases}$$

Then,

$$\mathbb{E}_H f(H) = \mathbb{E}_H \mathbb{E}_y f(H_{x,y})\,.$$

Proof. The proof is straightforward, and can be found in Lemma 19 of [41]. □

[2] We can chose the algorithm's unitaries between oracle calls to be all the same by introducing a "clock register" that keeps track of the number of oracle calls made so far.

The following is a technical lemma about a quantum adversary not being able to distinguish between samples from $H(U_\lambda)$ and from U_m, even when given oracle access to H, where the function $H : \{0,1\}^\lambda \to \{0,1\}^m$ is sampled uniformly at random.

Consider the following game between a challenger and a quantum adversary $\mathcal{A}$, specified by $\lambda, m \in \mathbb{N}$, and a distribution X over $\{0,1\}^\lambda$,

- The challenger samples a uniformly random function $H : \{0,1\}^\lambda \to \{0,1\}^m$ and $b \leftarrow \{0,1\}$.
- If $b = 0$: the challenger samples $x \leftarrow X$, sends $H(x)$ to $\mathcal{A}$.
  If $b = 1$: the challenger samples uniformly $z \leftarrow \{0,1\}^m$, sends z to $\mathcal{A}$.
- $\mathcal{A}$ additionally gets oracle access to H. $\mathcal{A}$ returns a bit b' to the challenger.

$\mathcal{A}$ wins if $b = b'$. Let $\mathsf{Dist}(\mathcal{A}, \lambda, m, X)$ be a random variable for the outcome of the game.

Lemma 3. *For any adversary $\mathcal{A}$ making q oracle queries, any family of distributions $\{X_\lambda : \lambda \in \mathbb{N}\}$ where for all λ, X_λ is a distribution over $\{0,1\}^\lambda$, for any polynomially bounded function $m : \mathbb{N} \to \mathbb{N}$, there exists a negligible function μ such that, for any $\lambda \in \mathbb{N}$, the following holds:*

$$\Pr[\mathsf{Dist}(\mathcal{A}, \lambda, m(\lambda), X_\lambda) = 1] \leq \frac{1}{2} + (3q+2)qM + \mu(\lambda)\,,$$

where M is a quantity that is negligible in λ if $2^{-\mathbf{H}_{\min}(X_\lambda)/2}$ is negligible in λ.

Corollary 1. *For any query-bounded adversary $\mathcal{A}$, any $\epsilon > 0$, any family of distributions $\{X_\lambda : \lambda \in \mathbb{N}\}$, where X_λ is a distribution over $\{0,1\}^\lambda$ with $\mathbf{H}_{\min}(X_\lambda) > \lambda^\epsilon$ for every λ, for any polynomially bounded function $m : \mathbb{N} \to \mathbb{N}$, there exists a negligible function μ such that, for any $\lambda \in \mathbb{N}$, the following holds:*

$$\Pr[\mathsf{Dist}(\mathcal{A}, \lambda, m(\lambda), X_\lambda) = 1] \leq \frac{1}{2} + \mu(\lambda)\,.$$

The key step in the proof of Lemma 3 is captured by the one-way-to-hiding lemma [127, 12][3]. We restate it here following our notation (and provide a proof a for completeness). Informally, the lemma gives an upper bound on an adversary's advantage (when given access to a uniformly random function $H : \{0,1\}^n \to \{0,1\}^m$) at distinguishing between a sample drawn from $H(U_n)$ and a sample drawn from U_m. The upper bound is in terms of the probability that the adversary queries the oracle at the pre-image of the sample at some point during its execution. Equivalently, given two oracles that are identical except on a single input (or more generally on a subset of the

[3] While additional improved variants of the one-way to hiding lemma were developed [31, 93], any of them suffices for our asymptotic analysis.

inputs), the advantage of an adversary at distinguishing the two oracles is bounded above in terms of the probability that the adversary queries at the differing point (or at a point in the subset where they differ) at some point during its execution.

Lemma 4. *Let $\lambda, m \in \mathbb{N}$. For any $q \in \mathbb{N}$, any unitaries U, any family of states $\{|\psi_x\rangle\}_{x\in\mathcal{X}}$, any complete pair of orthogonal projectors (Π^0, Π^1) and any distribution X on $\{0,1\}^\lambda$, it holds that:*

$$\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\|\Pi^0(UO^H)^q\left(|H(x)\rangle \otimes |\psi_x\rangle\right)\|^2 + \frac{1}{2}\mathbb{E}_H\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1(UO^H)^q\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 \leq \frac{1}{2} + (3q+2)qM\,, \tag{3.2}$$

where O^H is the oracle unitary for $H : \{0,1\}^\lambda \to \{0,1\}^m$, and M is given by

$$\begin{aligned} M =& \frac{1}{2}\,\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\mathbb{E}_k\|\,|x\rangle\langle x|\,(UO^{H_{x,z}})^k\,|z\rangle \otimes |\psi_x\rangle\,\| \\ &+\frac{1}{2}\,\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\mathbb{E}_k\|\,|x\rangle\langle x|\,(UO^{H})^k\,|z\rangle \otimes |\psi_x\rangle\,\|\,. \end{aligned} \tag{3.3}$$

Moreover, M is negligible if and only if the second term in M is negligible.

The lemma holds also when the states $|\psi_x\rangle$ are not necessarily pure (but we write them as pure states for ease of notation).

Proof. For any $x \in \{0,1\}^\lambda$, define $V_x^H = \left(UO^H(I - |x\rangle\langle x|)\right)^q$ and define $W_x^H = UO^H - V_x^H$. Then,

$$\begin{aligned} &\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\|\Pi^0(UO^H)^q\left(|H(x)\rangle \otimes |\psi_x\rangle\right)\|^2 + \frac{1}{2}\mathbb{E}_H\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1(UO^H)^q\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 \\ =&\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^0(UO^{H_{x,z}})^q\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 + \frac{1}{2}\mathbb{E}_H\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1(UO^H)^q\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 \\ =&\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^0(V_x^{H_{x,z}} + W_x^{H_{x,z}})\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 \\ &\quad + \frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1(V_x^H + W_x^H)\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 \\ \leq&\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^0V_x^{H_{x,z}}\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 + \frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1V_x^H\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 \\ &+\frac{1}{2}(3q+2)q\,\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\mathbb{E}_k\|\,|x\rangle\langle x|\,(UO^{H_{x,z}})^k\,|z\rangle \otimes |\psi_x\rangle\,\| \\ &+\frac{1}{2}(3q+2)q\,\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\mathbb{E}_k\|\,|x\rangle\langle x|\,(UO^{H})^k\,|z\rangle \otimes |\psi_x\rangle\,\| \\ =&\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^0V_x^{H_{x,z}}\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 + \frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1V_x^H\left(|z\rangle \otimes |\psi_x\rangle\right)\|^2 \\ &+(3q+2)q\,M \end{aligned} \tag{3.4}$$

where the first equality uses Lemma 2, and the inequality uses Lemma 18 in [41].

In order to prove the desired inequality, it is sufficient to show that

$$\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^0 V_x^{H_{x,z}}\left(|z\rangle\otimes|\psi_x\rangle\right)\|^2+\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1 V_x^{H}\left(|z\rangle\otimes|\psi_x\rangle\right)\|^2\leq\frac{1}{2}\,. \tag{3.5}$$

Notice that $V_x^{H_{x,z}}=V_x^H$, since V_x^H projects onto the subspace orthogonal to x before every query to H. This implies that the LHS simplifies as

$$\begin{aligned}&\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^0 V_x^{H_{x,z}}\left(|z\rangle\otimes|\psi_x\rangle\right)\|^2+\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1 V_x^{H}\left(|z\rangle\otimes|\psi_x\rangle\right)\|^2\\ =&\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^0 V_x^{H}\left(|z\rangle\otimes|\psi_x\rangle\right)\|^2+\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\in X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1 V_x^{H}\left(|z\rangle\otimes|\psi_x\rangle\right)\|^2\\ =&\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\in X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|V_x^{H}\left(|z\rangle\otimes|\psi_x\rangle\right)\|^2\ \leq\ \frac{1}{2}\,,\end{aligned} \tag{3.6}$$

where to get the third line, we used the fact that Π^0,Π^1 are a complete pair of orthogonal projectors, and to get the last line we exploited properties of the Euclidean norm.

Combining (3.4) and (3.6) gives the desired inequality.

With a little extra work, one can show that M is negligible if and only if

$$\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X}\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1 V_x^{H}\left(|z\rangle\otimes|\psi_x\rangle\right)\|^2$$

is negligible. We refer the reader to the proof of Theorem 3 in [12] for the full details. □

Proof of Lemma 3. Without loss of generality, let $\mathcal{A}$ be specified by a unitary U, the oracle unitary O^H and a measurement given by projectors Π_0 and $\Pi_1=\mathbb{1}-\Pi_0$, so that the unitary part of $\mathcal{A}$'s algorithm is $(UO^H)^q$, where q the number of oracle queries made by $\mathcal{A}$. Then, $\mathcal{A}$'s winning probability is precisely given by,

$$\begin{aligned}&\Pr[\mathsf{Dist}(\mathcal{A},\lambda,m,X_\lambda)=1]\\ &=\frac{1}{2}\mathbb{E}_H\mathbb{E}_{x\leftarrow X_\lambda}\|\Pi^0(UO^H)^q\,|H(x)\rangle\,\|^2+\frac{1}{2}\mathbb{E}_H\mathbb{E}_{z\leftarrow\{0,1\}^m}\|\Pi^1(UO^H)^q\,|z\rangle\,\|^2\,,\end{aligned} \tag{3.7}$$

where we omit writing ancilla qubits initialized in the zero state that $(UO^H)^q$ might be acting on.

Then, by Lemma 4, we have

$$\Pr[\mathsf{Dist}(\mathcal{A},\lambda,m,X_\lambda)=1]\leq\frac{1}{2}+(3q+2)q\,M \tag{3.8}$$

where M is the quantity given by

$$\begin{aligned}M=&\ \frac{1}{2}\,\mathbb{E}_H\mathbb{E}_{x\leftarrow X_\lambda}\mathbb{E}_{z\leftarrow\{0,1\}^m}\mathbb{E}_k\|\,|x\rangle\langle x|\,(UO^{H_{x,z}})^k\,|z\rangle\,\|\\ &+\frac{1}{2}\,\mathbb{E}_H\mathbb{E}_{x\leftarrow X_\lambda}\mathbb{E}_{z\leftarrow\{0,1\}^m}\mathbb{E}_k\|\,|x\rangle\langle x|\,(UO^{H})^k\,|z\rangle\,\|\,.\end{aligned} \tag{3.9}$$

Moreover, by Lemma 4, M is negligible if and only if the second term,

$$\mathbb{E}_H \mathbb{E}_{x \leftarrow X_\lambda} \mathbb{E}_{z \leftarrow \{0,1\}^m} \mathbb{E}_k \| \, |x\rangle \langle x| \, (UO^H)^k \, |z\rangle \, \|,$$

is negligible. Hence, it suffices to bound the above term. Notice that for any fixed H, z and k,

$$\begin{aligned} &\mathbb{E}_{x \leftarrow X_\lambda} \| \, |x\rangle \langle x| \, (UO^H)^k \, |z\rangle \, \| \\ &\quad \leq \sqrt{\mathbb{E}_{x \leftarrow X_\lambda} \| \, |x\rangle \langle x| \, (UO^H)^k \, |z\rangle \, \|^2} \\ &\quad \leq 2^{-\mathbf{H}_{\min}(X_\lambda)/2} \, , \end{aligned} \tag{3.10}$$

where the first inequality follows from Jensen's inequality (for concave functions), and the second inequality uses the fact that the state $(UO^H)^k \, |z\rangle$ does not depend on x, and hence the quantity under the square root is bounded above by the optimal probability of correctly predicting a sample from X, which is, by definition, $2^{-\mathbf{H}_{\min}(X)}$. Therefore, M is negligible so long as $2^{-\mathbf{H}_{\min}(X_\lambda)}$ is negligible. □

Finally, we define the notion of indistinguishability of ensembles of quantum states *in the QROM*. This is similar to Definition 2.

Definition 15 (Indistinguishability of ensembles of quantum states in the QROM)**.** *Let $m : \mathbb{N} \to \mathbb{N}$ and $p : \mathbb{N} \to \mathbb{N}$ be polynomially bounded functions, and let ϱ_λ^H and σ_λ^H be $p(\lambda)$-qubit states, for $H \in \mathsf{Bool}(\lambda, m(\lambda))$. We say that $\{\varrho_\lambda^H\}_{\lambda \in \mathbb{N}, H \in \mathsf{Bool}(\lambda, m(\lambda))}$ and $\{\sigma_\lambda^H\}_{\lambda \in \mathbb{N}, H \in \mathsf{Bool}(\lambda, m(\lambda))}$ are quantum computationally indistinguishable ensembles of quantum states, denoted by $\varrho_\lambda^H \approx_c \sigma_\lambda^H$, if, for any* QPT *distinguisher $\mathcal{D}^H$ 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 $\lambda \in \mathbb{N}$,*

$$\mathbb{E}_H \Big| \Pr[\mathcal{D}^H(\varrho_\lambda^H \otimes \nu_\lambda) = 1] - \Pr[\mathcal{D}^H(\sigma_\lambda^H \otimes \nu_\lambda) = 1] \Big| \leq \mathsf{negl}(\lambda) \, .$$

3.3 Quantum Copy-Protection

Our definition of a secure copy-protection scheme is essentially identical to the notion in [2]. We elaborate on the differences in Section 3.3.

Definition 16 (Quantum copy-protection scheme)**.** *Let $\mathscr{F} = \bigcup_{\lambda \in \mathbb{N}} \mathcal{F}_\lambda$ be a class of efficiently computable functions $f : \mathcal{X} \to \mathcal{Y}$ with domain $\mathcal{X}$ and range $\mathcal{Y}$. A quantum copy-protection (QCP) scheme for the class $\mathscr{F}$ is a pair of* QPT *algorithms* QCP = (Protect, Eval) *defined as follows:*

QCP.Protect$(1^\lambda, d_f) \to \varrho$: *takes as input the security parameter 1^λ and a classical description d_f of a function $f \in \mathcal{F}_\lambda$, and outputs a (possibly mixed) quantum state ϱ.*

$\mathsf{QCP.Eval}(1^\lambda, \varrho, x) \longrightarrow \varrho' \otimes |y\rangle\langle y|$: *takes as input the security parameter* 1^λ*, a quantum state* ϱ *and an input* $x \in \mathcal{X}$*, and outputs a bipartite state* $\varrho' \otimes |y\rangle\langle y|$ *with* $y \in \mathcal{Y}$*.*

Slightly abusing notation, we occasionally ignore the post-evaluation state ϱ' and simply identify the output of the procedure $\mathsf{QCP.Eval}(1^\lambda, \varrho, x)$ with a classical outcome denoted by $y \in \mathcal{Y}$.

We say that a QCP scheme is *ϵ-correct* if, for any $\lambda \in \mathbb{N}$, any $f \in \mathcal{F}_\lambda$, and any input $x \in \mathcal{X}$ to f:

$$\Pr\left[\mathsf{QCP.Eval}(1^\lambda, \varrho, x) = f(x) \ : \ \varrho \leftarrow \mathsf{QCP.Protect}(1^\lambda, d_f)\right] \geq 1 - \epsilon(\lambda).$$

Note that the probability above comes from the procedure $\mathsf{QCP.Eval}$ of the QCP scheme. If $\epsilon(\lambda) = \mathsf{negl}(\lambda)$, we simply call a copy-protection scheme *correct*. By the Gentle Measurement Lemma [134] it is easy to see that a ϵ-correct scheme is reusable in the following sense: after performing $\mathsf{QCP.Eval}$ to ϱ it is possible to rewind the procedure to obtain a state that is within trace distance $\sqrt{\epsilon}$ of the original state ϱ.

Informally, we say that a QCP scheme $\mathsf{QCP} = (\mathsf{Protect}, \mathsf{Eval})$ is *secure* if no QPT adversary can produce two "copies" of a copy-protected program $\varrho \leftarrow \mathsf{QCP.Protect}(1^\lambda, d_f)$ that can both be used to evaluate f. We formalise the security of copy-protection schemes by means of the following security experiment.

Definition 17 (Piracy experiment)**.** *Let* $\mathsf{QCP} = (\mathsf{Protect}, \mathsf{Eval})$ *be a copy-protection scheme for a class of functions* $\mathscr{F} = \bigcup_{\lambda\in\mathbb{N}} \mathcal{F}_\lambda$ *with domain* $\mathcal{X}$ *and range* $\mathcal{Y}$*. Let* $\mathcal{D}_{\mathscr{F}} = \{\mathcal{D}_{\mathcal{F}_\lambda}\}_{\lambda\in\mathbb{N}}$ *be an ensemble of distributions over* $\mathcal{F}_\lambda$ *and let* $\mathcal{D}_{\mathcal{X}} = \{\mathcal{D}_{\mathcal{X}}(f)\}_{f\in\mathcal{F}_\lambda}$ *be an ensemble of challenge distributions over function inputs* $\mathcal{X}$*. The security game (which we call the* piracy experiment*) takes place between a challenger and an adversary consisting of a triplet of* QPT *algorithms* $(\mathcal{A}, \mathcal{B}, \mathcal{C})$*:*

1. *The challenger samples* $f \leftarrow \mathcal{D}_{\mathcal{F}_\lambda}$ *and sends the program* $\varrho \leftarrow \mathsf{QCP.Protect}(1^\lambda, d_f)$ *to* $\mathcal{A}$*.*

2. $\mathcal{A}$ *applies an efficient* CPTP *map to map* ϱ *into a bipartite state* ϱ_{BC} *on systems* BC*, and sends system* B *to* $\mathcal{B}$ *and system* C *to* $\mathcal{C}$ *(who are not allowed to communicate from this step onward).*

3. *The challenger samples a pair* $(x_B, x_C) \leftarrow \mathcal{D}_{\mathcal{X}}(f) \times \mathcal{D}_{\mathcal{X}}(f)$*, and sends* x_B *to* $\mathcal{B}$ *and* x_C *to* $\mathcal{C}$*.*

4. $\mathcal{B}$ *and* $\mathcal{C}$ *output values* $y_B \in \mathcal{Y}$ *and* $y_C \in \mathcal{Y}$*, respectively, and send them to the challenger. The challenger outputs* 1*, if* $y_B = f(x_B)$ *and* $y_C = f(x_C)$ *(i.e., the adversary has succeeded) and* 0*, otherwise (i.e., the adversary has failed).*

We let the random variable $\mathsf{PiracyExp}^{\mathsf{QCP}}_{\mathcal{D}_{\mathscr{F}},\mathcal{D}_{\mathcal{X}}}\big(1^\lambda, (\mathcal{A}, \mathcal{B}, \mathcal{C})\big)$ *denote the output bit of the challenger.*

Definition 18 (Secure quantum copy-protection)**.** *Let* $\mathsf{QCP} = (\mathsf{Protect}, \mathsf{Eval})$ *be a QCP scheme for a class of functions* $\mathscr{F} = \bigcup_{\lambda \in \mathbb{N}} \mathcal{F}_\lambda$*. Let* $\mathcal{D}_{\mathscr{F}} = \{\mathcal{D}_{\mathcal{F}_\lambda}\}_{\lambda \in \mathbb{N}}$ *be an ensemble of distributions over* $\mathcal{F}_\lambda$ *and let* $\mathcal{D}_{\mathcal{X}} = \{\mathcal{D}_{\mathcal{X}}(f)\}_{f \in \mathcal{F}_\lambda}$ *be an ensemble of distributions over* $\mathcal{X}$*. Then,* $\mathsf{QCP} = (\mathsf{Protect}, \mathsf{Eval})$ *is called* $(\mathcal{D}_{\mathscr{F}}, \mathcal{D}_{\mathcal{X}})$*-secure if, for any triplet of* QPT *algorithms* $(\mathcal{A}, \mathcal{B}, C)$*, it holds that*

$$\Pr\left[\mathsf{PiracyExp}^{\mathsf{QCP}}_{\mathcal{D}_{\mathscr{F}}, \mathcal{D}_{\mathcal{X}}}\left(1^\lambda, (\mathcal{A}, \mathcal{B}, C)\right) = 1\right] \leq p^{\mathrm{triv}}_{\mathcal{D}_{\mathscr{F}}, \mathcal{D}_{\mathcal{X}}} + \mathsf{negl}(\lambda),$$

where $p^{\mathrm{triv}}_{\mathcal{D}_{\mathscr{F}}, \mathcal{D}_{\mathcal{X}}}$ *is the trivial winning probability that is always possible due to correctness:* $\mathcal{A}$ *forwards the original copy-protected program to one of the parties, say* $\mathcal{B}$ *(who then evaluates it to obtain the correct output), and the other party, say* C*, has to guess at random.*

Finally, we conclude this section a brief discussion regarding prior definitions of security.

Comparison with existing definitions of copy-protection

Our definition is very similar to the original security definition first proposed by Aaronson [2]. The only difference is the following. In [2], a scheme is δ-secure if for any bounded adversary who tries to create $k + 1$ programs upon receiving k copy-protected copies the average number of input challenges answered correctly is $k(1 + \delta)$. In contrast, in our definition we say that the scheme is secure if no adversary can succeed with non-negligible advantage beyond the trivial guessing probability. In our work, we exclusively focus on the case of $k = 1$.

3.4 Quantum Copy-Protection of Multi-Bit Point Functions

In this section, we make a conceptual connection between unclonable quantum encryption and quantum copy-protection. Our main result is a quantum copy-protection scheme for multi-bit point functions which we obtain from any unclonable encryption scheme with a so-called "wrong-key detection mechanism." Canetti et al. [46] previously introduced a similar property for classical encryption in the context of point function obfuscation.

A private-key quantum encryption of classical messages (QECM) scheme is a procedure that takes as input a key and a plaintext in the form of a classical bit string in a quantum register, and produces a ciphertext in the form of a quantum state. We formalise this notion in Definition 19.

Definition 19 (Quantum encryption scheme of classical messages)**.** *Let* $\lambda \in \mathbb{N}$ *be the security parameter. A quantum encryption of classical messages (QECM) scheme with key space* $\mathcal{K}$ *and plaintext space* $\mathcal{M}$ *is a triplet* $\mathsf{QECM} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ *consisting of* QPT *algorithms:*

- $\mathsf{KeyGen}(1^\lambda) \to k$*: takes as input the security parameter and outputs a key* $k \in \mathcal{K}$*.*
- $\mathsf{Enc}(k, m) \to \varrho$*: takes as input a key* k *and a message* $m \in \mathcal{M}$ *and outputs* $\varrho \in \mathcal{D}(\mathcal{H}_A)$*.*

- $\mathsf{Dec}(k, \varrho) \to \sigma$: *takes as input a key k and a quantum ciphertext $\varrho \in \mathcal{D}(\mathcal{H}_A)$, and outputs a plaintext in the form of a state $\sigma \in \mathcal{D}(\mathcal{H}_M)$, where $\mathcal{H}_M = \mathrm{span}\{|m\rangle : m \in \mathcal{M}\}$.*

We use the notation Enc_k for the map $m \mapsto \mathsf{Enc}(k, m)$, and likewise for Dec_k. Note that Enc_k can naturally be extended to quantum inputs $\sigma \in \mathcal{D}(\mathcal{H}_M)$ with $\mathcal{H}_M = \mathrm{span}\{|m\rangle : m \in \mathcal{M}\}$ via

$$\mathsf{Enc}_k(\sigma) = \sum_{m \in \mathcal{M}} \mathrm{Tr}[|m\rangle\langle m|\, \sigma]\, \mathsf{Enc}_k(|m\rangle\langle m|)\,.$$

A QECM scheme is called correct *if, for all $m \in \mathcal{M}$ and $k \in \mathrm{supp}\, \mathsf{KeyGen}(1^\lambda)$,*

$$\mathrm{Tr}|m\rangle\langle m|\, \mathsf{Dec}_k \circ \mathsf{Enc}_k(m) \geq 1 - \mathsf{negl}(\lambda).$$

Note that we typically assume that the key space is given by $\mathcal{K} = \{0, 1\}^\lambda$ and that the plaintext space $\mathcal{M}$ and ciphertext space consist of inputs of at most $\mathsf{poly}(\lambda)$ many bits and qubits, respectively.

We use the following notion of indistinguishable encryptions for general symmetric-key quantum encryption schemes introduced by Alagic et al. [10].

Definition 20 (Indistinguishable security)**.** *A QECM scheme $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ has* indistinguishable encryptions *(or is* IND*-secure) if, for every* QPT *adversary $\mathcal{A} = (\mathcal{M}_{\mathcal{A}}, \mathcal{D})$ consisting of an (adversarial) message sampling procedure $\mathcal{M}_{\mathcal{A}}$ and a distinguisher $\mathcal{D}$,*

$$\Big|\Pr\big[\mathcal{D}\big((\mathsf{Enc}_k \otimes \mathbb{1}_E)\varrho_{ME}\big) = 1\big] - \Pr\big[\mathcal{D}\big((\mathsf{Enc}_k \otimes \mathbb{1}_E)(|0\rangle\langle 0|_M \otimes \varrho_E)\big) = 1\big]\Big| \leq \mathsf{negl}(\lambda)\,,$$

where we assume that $k \leftarrow \mathsf{KeyGen}(1^\lambda)$, $\varrho_{ME} \leftarrow \mathcal{M}_{\mathcal{A}}(1^\lambda)$ with $\varrho_E = \mathrm{tr}_M[\varrho_{ME}]$, and where $|0\rangle\langle 0|_M$ is the all-0 string in the plaintext register M.

Informally, we say that a quantum encryption scheme is *unclonable* if no QPT adversary can produce two "copies" of a quantum ciphertext which can each be decrypted with access to the private key. Before we make the notion of unclonable ciphertexts more precise, let us first introduce the following definition of a cloning attack due to Broadbent and Lord [41].

Definition 21 (Cloning attack)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. A cloning attack $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ against a QECM scheme $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ consists of the following* QPT *algorithms (which are parameterised by λ)*

- *(cloning map) $\mathcal{A} : \mathcal{D}(\mathcal{H}_A) \to \mathcal{D}(\mathcal{H}_B \otimes \mathcal{H}_C)$*
- *(1st decoder) $\mathcal{B} : \mathcal{K} \times \mathcal{D}(\mathcal{H}_B) \to \mathcal{D}(\mathcal{H}_M)$*
- *(2nd decoder) $\mathcal{C} : \mathcal{K} \times \mathcal{D}(\mathcal{H}_C) \to \mathcal{D}(\mathcal{H}_M)$*

where $\mathcal{K}$ and $\mathcal{H}_A$ are defined by the quantum encryption scheme Σ, i.e., $\mathcal{K}$ is the set of keys and $\mathcal{H}_A$ is the ciphertext system.

We remark that we only consider *efficient* cloning attacks throughout this work. This is in contrast with the definition of Broadbent and Lord [41] who consider CPTP maps more generally.

Definition 22 (Cloning experiment)**.** *Let $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ be a QECM scheme and let $\lambda \in \mathbb{N}$ be a parameter. We define the following security game, called the* cloning experiment*, which takes place between a challenger and a* QPT *adversary who executes a cloning attack $(\mathcal{A}, \mathcal{B}, \mathcal{C})$:*

1. *The challenger samples $k \leftarrow \mathsf{KeyGen}(1^\lambda)$ and $m \overset{\$}{\leftarrow} \mathcal{M}$, and sends $\varrho_A \leftarrow \mathsf{Enc}_k(m)$ to $\mathcal{A}$.*
2. *$\mathcal{A}$ maps ϱ_A into a bipartite state ϱ_{BC} on systems BC, and sends ϱ_{BC} to the challenger together with descriptions of ensembles of efficient quantum algorithms $\{\mathcal{B}_\kappa\}_{\kappa\in\mathcal{K}}$ and $\{\mathcal{C}_\kappa\}_{\kappa\in\mathcal{K}}$.*
3. *The challenger runs $\mathcal{B}_k$ on system B and $\mathcal{C}_k$ on system C of ϱ_{BC}, measures the output states in the computational basis to obtain outcomes m_B and m_C, and outputs 1 if $m = m_B = m_C$, and 0 otherwise.*

We let the random variable $\mathsf{CloneExp}_\Sigma\big(1^\lambda, (\mathcal{A}, \mathcal{B}, \mathcal{C})\big)$ denote the output bit of the challenger.

Building on the cloning experiment in 22, we then define unclonable security as follows.

Definition 23 (Unclonable Security)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. We say that a QECM scheme $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ with message space $\{0,1\}^{n(\lambda)}$ is $t(\lambda)$-unclonable secure if, for all cloning attacks $(\mathcal{A}, \mathcal{B}, \mathcal{C})$, it holds that*

$$\Pr\left[\mathsf{CloneExp}_\Sigma\big(1^\lambda, (\mathcal{A}, \mathcal{B}, \mathcal{C})\big) = 1\right] \leq 2^{-n(\lambda)+t(\lambda)} + \mathsf{negl}(\lambda).$$

Quantum encryption with wrong-key detection

Let us first formalize the “wrong-key detection mechanism” for quantum encryption schemes.

Definition 24 (Wrong-Key Detection)**.** *Let $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ be a* SKQES*. We say that the scheme satisfies the wrong-key detection (WKD) property if, for every $k' \neq k \leftarrow \mathsf{KeyGen}(1^\lambda)$:*

$$\|\mathsf{Dec}_{k'} \circ \mathsf{Enc}_k - \langle|\perp\rangle\langle\perp|\rangle\|_\diamond \leq \mathsf{negl}(\lambda).$$

Here, $\langle|\perp\rangle\langle\perp|\rangle(\cdot) = |\perp\rangle\langle\perp| \operatorname{Tr}[\cdot]$.

Next, we give a simple transformation that achieves WKD in the QROM.

Construction 1 (Generic Transformation for WKD in the QROM)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter and $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ be a* QECM *scheme. Fix a function $H : \{0,1\}^\lambda \to \{0,1\}^\ell$. The* QECM $\Pi_H = (\mathsf{KeyGen}', \mathsf{Enc}', \mathsf{Dec}')$ *scheme is defined by the following* QPT *algorithms:*

- KeyGen'*: on input 1^λ, run* $\mathsf{KeyGen}(1^\lambda)$ *to output $k \in \mathcal{K}$, with $\mathcal{K} = \{0,1\}^\lambda$.*
- Enc'*: on input m, run* $\mathsf{Enc} : \mathcal{K} \times \mathcal{D}(\mathcal{H}_M) \to \mathcal{D}(\mathcal{H}_C)$ *and output* $(\mathsf{Enc}_k(|m\rangle\langle m|), H(k))$.
- Dec'*: on input (ϱ, c), first check whether $H(k) = c$. Output $|\bot\rangle\langle\bot|$, if false. Otherwise, run* $\mathsf{Dec} : \mathcal{K} \times \mathcal{D}(\mathcal{H}_C) \to \mathcal{D}(\mathcal{H}_M)$ *and output* $\mathsf{Dec}_k(\varrho)$.

Lemma 18. *Let Π be any $t(\lambda)$-unclonable secure* QECM *and let $H : \{0,1\}^\lambda \to \{0,1\}^\ell$ be a hash function, for $\ell = 2\lambda$. Then, Construction 1 yields an $t(\lambda)$-unclonable secure* QECM *scheme Π_H with WKD in the QROM.*

Proof. Correctness is clearly preserved. Let us first verify the WKD property of the scheme $\Pi_H = (\mathsf{KeyGen}', \mathsf{Enc}', \mathsf{Dec}')$ in the QROM. Let $k \leftarrow \mathsf{KeyGen}'(1^\lambda)$. It is not hard to see that the WKD property depends on the collision probability for the event that $H(k) = H(k')$, for some $k' \in \{0,1\}^\lambda \setminus \{k\}$. In fact, we can express the quantum channel $\mathsf{Dec}'_{k'} \circ \mathsf{Enc}'_k$ as follows:

$$\mathsf{Dec}'_{k'} \circ \mathsf{Enc}'_k = \Pr[\textsc{Coll}]\mathsf{Dec}_{k'} \circ \mathsf{Enc}_k + (1 - \Pr[\textsc{Coll}])\langle|\bot\rangle\langle\bot|\rangle.$$

Moreover, by the birthday bound, we have

$$\begin{aligned}\Pr[\textsc{Coll}] &= \Pr_H[\exists k' \in \{0,1\}^\lambda \setminus \{k\} \,:\, H(k) = H(k')]\\ &\leq \sum_{k' \in \{0,1\}^\lambda \setminus \{k\}} \Pr_H[H(k) = H(k')] = \frac{2^\lambda - 1}{2^{2\lambda}}.\end{aligned}$$

Hence, we can readily verify the WKD property as follows:

$$\begin{aligned}&\|\mathsf{Dec}'_{k'} \circ \mathsf{Enc}'_k - \langle|\bot\rangle\langle\bot|\rangle\|_\diamond\\ &= \max_{\varrho_{MM'}} \|(\mathsf{Dec}'_{k'} \circ \mathsf{Enc}'_k \otimes \mathbb{1}_{M'})(\varrho_{MM'}) - (\langle|\bot\rangle\langle\bot|\rangle \otimes \mathbb{1}_{M'})(\varrho_{MM'})\|_1\\ &\leq \max_{\varrho_{MM'}} \|\Pr[\textsc{Coll}]\mathsf{Dec}_{k'} \circ \mathsf{Enc}_k(\varrho_{MM'}) - \Pr[\textsc{Coll}](\langle|\bot\rangle\langle\bot|\rangle \otimes \mathbb{1}_{M'})(\varrho_{MM'})\|_1\\ &\leq \Pr[\textsc{Coll}] \max_{\varrho_{MM'}}\Big(\|\varrho_{MM'}\|_1 + \|(\langle|\bot\rangle\langle\bot|\rangle \otimes \mathbb{1}_{M'})(\varrho_{MM'})\|_1\Big) \leq \frac{2^\lambda - 1}{2^{2\lambda-1}} = \mathsf{negl}(\lambda).\end{aligned}$$

For security, let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ and recall that $\mathsf{Enc}'_k(|m\rangle\langle m|) = (\mathsf{Enc}_k(|m\rangle\langle m|), H(k))$ according to Constr. 1. Let $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ be an adversary against the unclonable security game with respect to Π_H. We give a reduction from the unclonable security of Π. Suppose that $\mathcal{A}$ receives access to a re-programmed oracle $H_{k,z}$, where $H(k) = z$, and that $\mathcal{A}$ makes at most $q = \mathsf{poly}(\lambda)$ queries in total. Without loss of generality, we assume that $\mathcal{A}$ is specified by $(UO^H)^q$, for some unitary U. It suffices to argue that the following global states are negligibly close in trace distance:

$$\begin{aligned}&\mathbb{E}_H\mathbb{E}_k\mathbb{E}_z \,|H\rangle\langle H| \otimes |k\rangle\langle k| \otimes \left((UO^{H_{k,z}})^q \mathsf{Enc}_k(|m\rangle\langle m|) \otimes |z\rangle\langle z| \left((UO^{H_{k,z}})^q\right)^\dagger\right)\\ \approx\, &\mathbb{E}_H\mathbb{E}_k\mathbb{E}_z \,|H\rangle\langle H|\otimes|k\rangle\langle k|\otimes\left((UO^H)^q \mathsf{Enc}_k(|m\rangle\langle m|)\otimes|z\rangle\langle z|\left((UO^H)^q\right)^\dagger\right).\end{aligned} \tag{3.11}$$

We use the one-way-to-hiding (Lemma 4) to deduce that the above distance is negligible, so long as the following quantity is negligible:

$$\mathbb{E}_H\mathbb{E}_k\mathbb{E}_z\mathbb{E}_v \mathrm{Tr}|k\rangle\langle k|\,(UO^H)^v\mathsf{Enc}_k(|m\rangle\langle m|)\otimes|z\rangle\langle z|\,((UO^H)^v)^\dagger. \tag{3.12}$$

Suppose for the sake of contradiction that the latter is non-negligible. Then, we can construct an adversary that wins at the unlconable security game against Π. The reduction is straightforward: the adversary for the unlconable security game runs $\mathcal{A}$ (by simulating the random oracle H) to extract k. The adversary then decrypts the challenge ciphertext using k, and forwards the appropriate message m to the decoders $\mathcal{B}$ and C. By the assumption that the adversary succeeds with non-negligible probability, so does the adversary against the unclonable security of Π, yielding a contradiction.

Using Eq. 3.11, Lemma 6 and that Π is $t(\lambda)$-unclonable secure it follows that, for all QPT cloning attacks $\mathcal{A} = (\mathcal{A}, \mathcal{B}, C)$ against Π_H, there exists a negligible function $\mu(\lambda)$ such that:

$$\mathop{\mathbb{E}}_{m}\mathop{\mathbb{E}}_{k\leftarrow\mathcal{K}} \mathrm{Tr}(|m\rangle\langle m| \otimes |m\rangle\langle m|)(\mathcal{B}_k \otimes C_k) \circ \mathcal{A} \circ \mathsf{Enc}_k(|m\rangle\langle m|) \leq 2^{-\lambda+t(\lambda)} + \mu(\lambda).$$

We conclude that Π_H is $t(\lambda)$-unclonable secure. □ □

Quantum copy-protection of multi-bit point functions from unclonable encryption schemes with wrong-key detection

We are now ready to state our quantum copy-protection scheme for multi-bit point functions which we obtain from any unclonable encryption scheme with the aforementioned "wrong-key detection mechanism." Here, we consider multi-bit point functions $P_{y,m}$ of the form

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

where $y, m \in \{0,1\}^\lambda$. Our construction is the following:

Construction 2 (Quantum copy-protection scheme for multi-bit point functions)**.**
To construct a QCP scheme for multi-bit point functions with input and output sizes λ, respectively, let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ *be a* QECM *with WKD, with security parameter and message length equal to λ. We define the QCP scheme* QCP = (Protect, Eval) *as follows:*

- $\mathsf{QCP.Protect}(1^\lambda, P_{y,m})$: *Takes as input a security parameter λ and a multi-bit point function $P_{y,m}$, succinctly specified by the marked input y (of size λ) and message m (of size λ), and outputs the quantum ciphertext given by $\mathsf{Enc}_y(m)$.*
- $\mathsf{QCP.Eval}(1^\lambda, \varrho; x)$: *Takes as input a security parameter λ, an alleged copy-protected program ϱ, and a string $x \in \{0,1\}^\lambda$ (the input on which the program is to be evaluated). Appends an ancillary qubit in the $|0\rangle$ state. Then, coherently[4] performs a two-outcome measurement to check whether $\mathsf{Dec}_x(\varrho)$ is in the state $|\perp\rangle\langle\perp|$, or not, and stores the resulting bit in the ancilla. If true, output 0^λ. Otherwise, rewind the procedure and measure in the standard basis to obtain a message m'.*

Before stating our main theorem on the security of Construction 2, we define the following two classes of distributions with respect to multi-bit point functions $P_{y,m}$ of the form

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

First, we let $\mathcal{D} = \{D_\lambda\}$ be an ensemble of distributions D_λ over multi-bit point functions $P_{y,m}$ over $\{0,1\}^\lambda$ that sample a marked input y as well an output message m uniformly at random with respect to $\{0,1\}^\lambda$. Further, by $\mathcal{D}' = \{D_y\}$ we denote an arbitrary ensemble of challenge distributions, where each D_y is a distribution of challenge input pairs to the program.

We prove the following theorem on the security of Construction 2:

Theorem 3. *Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ be any $t(\lambda)$-unclonable secure* QECM *with WKD such that $\lambda - t = \omega(\log \lambda)$. Then, Construction 2 yields a secure quantum copy-protection scheme* QCP *for multi-bit point functions with respect to the pair of ensembles $(\mathcal{D}, \mathcal{D}')$, against computationally-bounded adversaries.*

Proof. The correctness of the QCP scheme $\mathsf{QCP} = (\mathsf{Protect}, \mathsf{Eval})$ follows directly from the WKD property of the QECM. Let $\mathsf{Adv} = (\mathcal{A}, \mathcal{B}, C)$ denote the adversary for $\mathsf{PiracyExp}^{\mathsf{QCP}}_{D_\lambda, D_y}$. We consider two cases, namely when $p^{\text{triv}}_{D_\lambda, D_y} = 1$ and when $p^{\text{triv}}_{D_\lambda, D_y} < 1$ (depending on the challenge distribution $\mathcal{D}' = \{D_y\}$). In the former case, the scheme is trivially secure by definition and we are done. Hence, we will assume that $p^{\text{triv}}_{D_\lambda, D_y} < 1$ for the remainder of the proof. Note that, in this case, the distribution D_y has non-zero weight on the marked input y.

[4] If Dec is not unitary, performing this measurement coherently requires purifying it.

Let $(x_1, x_2) \leftarrow D_y$ denote the inputs received by the freeloaders $\mathcal{B}$ and $\mathcal{C}$ during the challenge phase. We can express the probability that Adv succeeds at $\mathsf{PiracyExp}^{\mathsf{QCP}}_{D_\lambda, D_y}$ as follows:

$$\begin{aligned}
&\Pr[\mathsf{Adv}\text{ wins}]\\
&=\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1\neq y\neq x_2]\cdot\Pr[x_1\neq y\neq x_2]+\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1=y\neq x_2]\cdot\Pr[x_1=y\neq x_2]\\
&+\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1\neq y=x_2]\cdot\Pr[x_1\neq y=x_2]+\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1=y=x_2]\cdot\Pr[x_1=y=x_2].
\end{aligned} \tag{3.13}$$

Without loss of generality, we assume that $\Pr[x_1 = y\neq x_2] \leq \Pr[x_1\neq y = x_2]$. Hence,

$$\begin{aligned}
\Pr[\mathsf{Adv}\text{ wins}] &\leq \Pr[\mathsf{Adv}\text{ wins}\,|\,x_1\neq y\neq x_2]\cdot\Pr[x_1\neq y\neq x_2]\\
&+\big(\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 = y\neq x_2]+\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1\neq y = x_2]\big)\cdot\Pr[x_1\neq y = x_2]\\
&+\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 = y = x_2]\cdot\Pr[x_1 = y = x_2].
\end{aligned} \tag{3.14}$$

Let us now state the following simple inequality. By first applying the union bound and then using that $\mathcal{B}$ and $\mathcal{C}$ are non-signalling, we find that:

$$\begin{aligned}
&\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 = y = x_2]\\
&=\Pr[\mathcal{B}\text{ succeeds}\wedge\mathcal{C}\text{ succeeds}\,|\,x_1 = y = x_2]\\
&\geq\Pr[\mathcal{B}\text{ succeeds}\,|\,x_1 = y = x_2]+\Pr[\mathcal{C}\text{ succeeds}\,|\,x_1 = y = x_2]-1\\
&=\Pr[\mathcal{B}\text{ succeeds}\,|\,x_1 = y \neq x_2]+\Pr[\mathcal{C}\text{ succeeds}\,|\,x_1 \neq y = x_2]-1\\
&\geq\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 = y \neq x_2]+\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 \neq y = x_2]-1.
\end{aligned} \tag{3.15}$$

Plugging this into Eq. (3.14), we obtain the following upper bound:

$$\begin{aligned}
\Pr[\mathsf{Adv}\text{ wins}] &\leq \Pr[\mathsf{Adv}\text{ wins}\,|\,x_1\neq y\neq x_2]\cdot\Pr[x_1\neq y\neq x_2]\\
&+\big(1+\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 = y = x_2]\big)\cdot\Pr[x_1\neq y = x_2]\\
&+\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 = y = x_2]\cdot\Pr[x_1 = y = x_2]\\
&\leq \Pr[x_1\neq y\neq x_2]+\Pr[x_1\neq y = x_2]+2\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 = y = x_2]\\
&= p^{\text{triv}}_{D_\lambda, D_y}+2\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 = y = x_2].
\end{aligned} \tag{3.16}$$

In the last line, we used the assumption that $\Pr[x_1 = y \neq x_2] \leq \Pr[x_1 \neq y = x_2]$ together with the following simple identity for the trivial guessing probability:

$$p^{\text{triv}}_{D_\lambda, D_y} = \Pr[x_1\neq y\neq x_2]+\max\big\{\Pr[x_1\neq y = x_2], \Pr[x_1 = y\neq x_2]\big\}.$$

We complete the proof by showing that $\Pr[\mathsf{Adv}\text{ wins}\,|\,x_1 = y = x_2] \leq \mathsf{negl}(\lambda)$. This implies that

$$\Pr[\mathsf{Adv}\text{ wins}] \leq p^{\text{triv}}_{D_\lambda, D_y}+\mathsf{negl}(\lambda). \tag{3.17}$$

Suppose that $\mathsf{Adv} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ succeeds with probability $p = \Pr[\mathsf{Adv}\text{ wins} \,|\, x_1 = y = x_2]$ on the challenge pair consisting of $x_1 = y$ and $x_2 = y$. We will use Adv to construct an adversary against the unclonable security of the QECM scheme Π. Consider the QPT adversary $\mathsf{Adv}' = (\mathcal{A}', \mathcal{B}', \mathcal{C}')$ against Π, which we define as follows:

- $\mathcal{A}'$ receives the state $\varrho = \mathsf{Enc}_y(|m\rangle\langle m|)$ and runs the pirate $\mathcal{A}$ on ϱ. Next, $\mathcal{A}'$ passes the two registers output by $\mathcal{A}$ to the decoders $\mathcal{B}'$ and $\mathcal{C}'$.
- The decoders $\mathcal{B}'$ and $\mathcal{C}'$ each receive the marked input y and then run the freeloaders $\mathcal{B}$ and $\mathcal{C}$, respectively, on the two registers prepared by $\mathcal{A}'$. Finally, the decoders output the outcomes obtained from running the freeloaders.

Since Π is $t(\lambda)$-unclonable secure, there exists a negligible $\mu(\lambda)$ such that:

$$\mathop{\mathbb{E}}_{m}\mathop{\mathbb{E}}_{y} \operatorname{Tr}(|m\rangle\langle m| \otimes |m\rangle\langle m|)(\mathcal{B}'_y \otimes \mathcal{C}'_y) \circ \mathcal{A}' \circ \mathsf{Enc}_k(|m\rangle\langle m|) \leq 2^{-\lambda+t(\lambda)} + \mu(\lambda).$$

Since $\lambda - t = \omega(\log \lambda)$, we conclude that Adv' succeeds with probability $p \leq \mathsf{negl}(\lambda)$. This completes the proof of Eq. (3.17), and thus the proof of Thm. 3. □

Finally, when applying the WKD transformation from Construction 1 to the $\log_2(9)$-unclonable encryption scheme by Broadbent and Lord [41], we obtain the following theorem:

Theorem 4. *There exists a* $\log_2(9)$*-unclonable secure* QECM *scheme with WKD for which Construction 2 yields a secure QCP scheme for multi-bit point functions with respect to the pair of ensembles* $(\mathcal{D}, \mathcal{D}')$*, against query-bounded (computationally bounded) adversaries in the QROM.*

3.5 Secure Software Leasing

In this section, consider a weaker notion of quantum copy-protection called "secure software leasing" (SSL) which was introduced in [17]. The crucial difference between the two notions lies in the fact that the scheme comes with a prescribed verification routine.

The syntax of a secure software leasing scheme is as follows.

Definition 25 (Secure software leasing)**.** *Let* $\mathscr{F} = \bigcup_{\lambda\in\mathbb{N}} \mathcal{F}_\lambda$ *be a class of efficiently computable functions* $f : \mathcal{X} \rightarrow \mathcal{Y}$ *with domain* $\mathcal{X}$ *and range* $\mathcal{Y}$*. A secure software leasing* (SSL) *scheme for* $\mathscr{F}$ *consists of* QPT *algorithms* $\mathsf{SSL} = (\mathsf{Gen}, \mathsf{Lease}, \mathsf{Eval}, \mathsf{Verify})$ *defined as follows:*

- $\mathsf{SSL.Gen}(1^\lambda)$ *takes as input the security parameter* λ *and outputs a secret key* sk.
- $\mathsf{SSL.Lease}(\mathsf{sk}, f)$ *takes as input a secret key* sk *and a function* $f \in \mathcal{F}_\lambda$*, and outputs a quantum state* ϱ_f*.*

- $\mathsf{SSL.Eval}(x, \varrho_f)$ *takes a string x as input to f together with a state ϱ_f, and outputs y' and a post-evaluation state $\tilde{\varrho}_f$.*
- $\mathsf{SSL.Verify}(\mathsf{sk}, f, \sigma)$ *takes as input the secret key* sk*, the function $f \in \mathcal{F}_\lambda$ and a state σ, and outputs* 1*, if σ is a valid lease state for f, and* 0 *otherwise.*

There exists a negligible function μ such that the scheme satisfies:

- *Correctness of evaluation: for all $\lambda \in \mathbb{N}$, for all $f \in \mathcal{F}_\lambda$, and for all x in the domain of f,*

$$\Pr\left[\mathsf{SSL.Eval}(x, \varrho) = f(x) : \varrho \leftarrow \mathsf{SSL.Lease}(\mathsf{sk}, f),\ \mathsf{sk} \leftarrow \mathsf{SSL.Gen}(1^\lambda)\right] \geq 1 - \mu(\lambda).$$

- *Correctness of verification: for all $\lambda \in \mathbb{N}$ and for all $f \in \mathcal{F}_\lambda$,*

$$\Pr\left[\mathsf{SSL.Verify}(\mathsf{sk}, f, \varrho) = 1 : \varrho \leftarrow \mathsf{SSL.Lease}(\mathsf{sk}, f),\ \mathsf{sk} \leftarrow \mathsf{SSL.Gen}(1^\lambda)\right] \geq 1 - \mu(\lambda).$$

Security is defined in terms of a security game between a lessor and an adversary $\mathcal{A}$ (the lessee). Informally, any secure software leasing (SSL) scheme should satisfy the following key property. After receiving a leased copy of f denoted by ϱ_f (generated using $\mathsf{SSL.Lease}$), the adversary should not be able to produce a quantum state σ on registers R_1 and R_2 such that:

- $\mathsf{SSL.Verify}$ deems the contents of register R_1 of $\sigma_{\mathsf{R}_1\mathsf{R}_2}$ to be valid, once it is returned.
- The adversary can still compute f (on inputs chosen by the lessor) from the post-measurement state in register R_2 given by $\sigma^*_{\mathsf{R}_2} \propto \mathrm{Tr}_{\mathsf{R}_1}\left[\Pi_1\left[(\mathsf{SSL.Verify}(\cdot)_{\mathsf{R}_1} \otimes \mathbb{1}_{\mathsf{R}_2})\sigma_{\mathsf{R}_1\mathsf{R}_2}\right]\right]$.

As in the case of quantum copy-protection schemes, we consider a *program ensemble* distribution $\mathcal{D} = \{D_{\mathcal{F}_\lambda}\}_{\lambda\in\mathbb{N}}$ and an *input challenge ensemble* of distributions $\{\mathcal{D}_{\mathcal{X}}(f)\}_{f\in\mathcal{F}_\lambda}$. To formalize the security of SSL schemes, we consider the following experiment.

Definition 26 (Piracy experiment for secure software leasing)**.** *Let* $\mathsf{SSL} = (\mathsf{Gen}, \mathsf{Lease}, \mathsf{Eval}, \mathsf{Verify})$ *be a secure software leasing scheme for a class of functions $\mathscr{F} = \bigcup_{\lambda\in\mathbb{N}} \mathcal{F}_\lambda$ with domain $\mathcal{X}$ and range $\mathcal{Y}$. Let $\mathcal{D}_{\mathscr{F}} = \{\mathcal{D}_{\mathcal{F}_\lambda}\}_{\lambda\in\mathbb{N}}$ be an ensemble of distributions over $\mathcal{F}_\lambda$ and let $\mathcal{D}_{\mathcal{X}} = \{\mathcal{D}_{\mathcal{X}}(f)\}_{f\in\mathcal{F}_\lambda}$ be an ensemble of challenge distributions over function inputs $\mathcal{X}$. The security game (which we call* piracy experiment*) takes place as follows between a lessor and a* QPT *adversary $\mathcal{A}$:*

1. *The lessor samples a function $f \leftarrow \mathcal{F}_\lambda$ and runs* $\mathsf{sk} \leftarrow \mathsf{SSL.Gen}(1^\lambda)$*. Then, the lessor runs* $\varrho \leftarrow \mathsf{SSL.Lease}(\mathsf{sk}, f)$*. The lessor sends ϱ to $\mathcal{A}$.*

2. *$\mathcal{A}$ outputs a (possibly entangled) state σ on two registers R_1 and R_2, and then sends the first register R_1 to the lessor.*

3. *For verification, the lessor runs* $\mathsf{SSL.Verify}$ *on input the secret key* sk*, a function $f \in \mathcal{F}_\lambda$ and the register R_1 of the state $\sigma_{\mathsf{R}_1\mathsf{R}_2}$. If* $\mathsf{SSL.Verify}$ *accepts, the lessor outputs* $\mathsf{ok} = 1$ *and lets the game continue, otherwise, the lessor outputs* $\mathsf{ok} = 0$ *and $\mathcal{A}$ loses.*

4. *The lessor samples $x \leftarrow \mathcal{D}_{\mathcal{X}}$, and sends x to the adversary $\mathcal{A}$.*

5. *$\mathcal{A}$ responds with a bit b. If $b = f(x)$, the lessor outputs* 1*. Otherwise, the lessor outputs* 0.

We let the random variable $\mathsf{PiracyExp}^{\mathsf{SSL}}_{\mathcal{D}_{\mathscr{F}},\mathcal{D}_{\mathcal{X}}}(1^\lambda, \mathcal{A})$ *denote the output bit of the challenger.*

We now give a formal definition of security.

Definition 27 (Security)**.** *Let* $\mathsf{SSL} = (\mathsf{Gen}, \mathsf{Lease}, \mathsf{Eval}, \mathsf{Verify})$ *be an SSL scheme for a class of functions $\mathscr{F} = \bigcup_{\lambda\in\mathbb{N}} \mathcal{F}_\lambda$. Let $\mathcal{D}_{\mathscr{F}} = \{\mathcal{D}_{\mathcal{F}_\lambda}\}_{\lambda\in\mathbb{N}}$ be an ensemble of distributions over $\mathcal{F}_\lambda$ and let $\mathcal{D}_{\mathcal{X}} = \{\mathcal{D}_{\mathcal{X}}(f)\}_{f\in\mathcal{F}_\lambda}$ be an ensemble of distributions over $\mathcal{X}$. Then,* SSL *is called $(\mathcal{D}_{\mathscr{F}}, \mathcal{D}_{\mathcal{X}})$-secure if, for any* QPT *algorithm $\mathcal{A}$, it holds that*

$$\Pr\left[\mathsf{PiracyExp}^{\mathsf{SSL}}_{\mathcal{D}_{\mathscr{F}},\mathcal{D}_{\mathcal{X}}}(1^\lambda, \mathcal{A}) = 1\right] \leq p^{\mathsf{triv},\mathsf{SSL}}_{\mathcal{D}_{\mathscr{F}},\mathcal{D}_{\mathcal{X}}} + \mathsf{negl}(\lambda),$$

Here, $p^{\mathsf{triv},\mathsf{SSL}}_{\mathcal{D}_{\mathscr{F}},\mathcal{D}_{\mathcal{X}}}$ is the trivial winning probability which corresponds to the guessing probability of the challenge distribution $\mathcal{D}_{\mathcal{X}}$. In other words,

$$p^{\mathsf{triv},\mathsf{SSL}}_{\mathcal{D}_{\mathscr{F}},\mathcal{D}_{\mathcal{X}}} = \max_{x\in\mathcal{X}} \mathbb{E}_{f\leftarrow\mathcal{D}_{\mathcal{F}_\lambda}} \hat{D}_f(x),$$

where $\hat{D}_f(x)$ is the probability that the correct answer to a challenge sampled from $\mathcal{D}_{\mathcal{F}_\lambda}$ is x.

3.6 Secure Software Leasing for Compute-and-Compare Programs

In this section, we show how to obtain an SSL scheme for a general class of compute-and-compare programs [131, 76]. A compute-and-compare program $\mathsf{CC}[f, y]$ is specified by an efficiently computable function $f : \{0,1\}^n \to \{0,1\}^m$ and a string $y \in \{0,1\}^m$ in its range, where

$$\mathsf{CC}[f, y](x) = \begin{cases} 1 & \text{if } f(x) = y\,, \\ 0 & \text{if } f(x) \neq y\,. \end{cases}$$

Note that point functions are a special case of compute-and-compare programs where the function f is the identity map.

In this section, we show how to lease $\mathsf{CC}[f, y]$ in the following simple way: the encoded program consists of (a description of) a function f in the clear, together with a quantumly encoded version of the point function with marked input y. In fact, it is straightforward to show that the SSL security of the extended construction reduces to the SSL security of the original point function scheme.

Secure software leasing for single-bit point functions

First, we show how to obtain an SSL scheme for single-bit point functions $\{P_y\}_{y\in\{0,1\}^\lambda}$ of the form

$$P_y(x) = \begin{cases} 1 & \text{if } x = y\,, \\ 0 & \text{if } x \neq y\,. \end{cases}$$

Specifically, we will focus on our attention on *unpredictable point function distributions* consisting of a distribution D_λ over point functions on $\{0,1\}^\lambda$ such that $P_y \leftarrow D_\lambda$ satisfies $\mathbf{H}_{\min}(y) \geq \lambda^\epsilon$ for some $\epsilon > 0$. We will now state our SSL scheme for single-bit point functions. In the following, we omit the procedure $\mathsf{SSL.Gen}$ as we do not require it in our construction.

Construction 3 (SSL scheme for point functions)**.** *Let λ be the security parameter, and let $H : \{0,1\}^{m(\lambda)} \to \{0,1\}^\lambda$ and $G : \{0,1\}^n \to \{0,1\}^{m(\lambda)}$ be hash functions, where $m(\lambda) \geq \lambda$. Consider the following secure software leasing* (SSL) *scheme* ($\mathsf{SSL.Lease}, \mathsf{SSL.Eval}, \mathsf{SSL.Verify}$) *for point functions P_y with marked input $y \in \{0,1\}^n$:*

- $\mathsf{SSL.Lease}(1^\lambda, y)$*: Takes as input a security parameter λ and a point function P_y, succinctly specified by the marked input y (of size n)*
  - *Set $\theta = G(y)$.*
  - *Sample $v \leftarrow \{0,1\}^{m(\lambda)}$ uniformly at random and let $z = H(v)$.*
  - *Output $(|v^\theta\rangle\,, z)$.*
- $\mathsf{SSL.Eval}(1^\lambda, (\varrho, z); x)$*: Takes as input a security parameter λ, a program (ϱ, z), and a string $x \in \{0,1\}^n$ (the input on which the program is to be evaluated).*
  - *Set $\theta' = G(x)$.*
  - *Apply Hadamards $H^{\theta'} = H^{\theta'_1} \otimes \cdots \otimes H^{\theta'_\lambda}$ to ϱ. Append $n+1$ ancillary qubits, all in state $|0\rangle$, and compute the hash function H with input ϱ into the first n of them (possibly making use of additional ancillary qubits). Then, coherently measure whether the first n ancilla qubits are in state $|z\rangle$, recording the result in the last ancilla qubit, uncompute the hash function H and undo the Hadamards $H^{\theta'}$. Finally, measure the last ancilla qubit to obtain a bit b as output.*
- $\mathsf{SSL.Verify}(1^\lambda, y, z, \sigma)$*: Apply H^θ to the input state σ, where $\theta = G(y)$, and measure in the standard basis. Output* 1 *if the result is v such that $H(v) = z$, and* 0 *otherwise.*

The correctness property of Construction 3 according to Definition 25 is immediate to verify. Before stating our main theorem on the security of Construction 3, we introduce a few classes of distributions over point functions and input challenges.

- $\mathcal{D}_{\mathsf{PF\text{-}UNP}}$. The class of *unpredictable point function distributions* $\mathcal{D}_{\mathsf{PF\text{-}UNP}}$ consists of ensembles $D = \{D_\lambda\}$ where D_λ is a distribution over point functions on $\{0,1\}^\lambda$ such that $P_y \leftarrow D_\lambda$ satisfies $\mathbf{H}_{\min}(y) \geq \lambda^\epsilon$ for some $\epsilon > 0$.

We also define the following class of distributions over input challenges.

- $\mathcal{D}_{\mathsf{PF\text{-}Chall\text{-}SSL}}$. An ensemble $D = \{D_y\}$, where each D_y is a distribution over $\{0,1\}^{|y|}$, belongs to the class $\mathcal{D}_{\mathsf{PF\text{-}Chall\text{-}SSL}}$ if there exists an efficiently sampleable family $\{X_\lambda\}$ of distributions over $\{0,1\}^\lambda$ with $\mathbf{H}_{\min}(X_\lambda) \geq \lambda^\epsilon$, for some $\epsilon > 0$, such that D_y is the following distribution (where $\lambda = |y|$):

  – with probability $1/2$, output y.

  – with probability $1/2$, sample $x \leftarrow X_\lambda$, and output x.

  We say the ensemble D is *specified* by the ensemble X_λ.

We finally define two classes of distributions over pairs of programs and challenges.

- $\mathcal{D}_{\mathsf{PF\text{-}pairs\text{-}stat\text{-}SSL}}$. This consists of pairs of ensembles $\left(D = \{D_\lambda\}, D' = \{D'_y\}\right)$ where $D \in \mathcal{D}_{\mathsf{PF\text{-}UNP}}$ and $D' \in \mathcal{D}_{\mathsf{PF\text{-}Chall\text{-}SSL}}$ satisfying the following. Let D' be parametrized by the family $\{X_\lambda\}$ (following the notation introduced above), and denote by $\mathsf{MarkedInput}(D_\lambda)$ the distribution over marked points in $\{0,1\}^\lambda$ induced by D_λ. Then, the families $\{X_\lambda\}$ and $\{\mathsf{MarkedInput}(D_\lambda)\}$ are statistically indistinguishable.

- $\mathcal{D}_{\mathsf{PF\text{-}pairs\text{-}comp\text{-}SSL}}$. This is defined in the same way as $\mathcal{D}_{\mathsf{PF\text{-}pairs\text{-}stat\text{-}SSL}}$, except that we only require $\{X_\lambda\}$ and $\{\mathsf{MarkedInput}(D_\lambda)\}$ to be *computationally* indistinguishable.

The following is our main result on the security of Construction 3.

Theorem 5. *The scheme of Construction 3, with $m(\lambda) = \mathsf{poly}(\lambda)$, is a secure software leasing scheme for point functions with respect to any pair of ensembles $(\mathcal{D}, \mathcal{D}') \in \mathcal{D}_{\mathsf{PF\text{-}pairs\text{-}stat\text{-}SSL}}$ ($\in \mathcal{D}_{\mathsf{PF\text{-}pairs\text{-}comp\text{-}SSL}}$), against query-bounded (computationally bounded) adversaries in the quantum random oracle model.*

Theorem 5 implies that, once a leased copy is successfully returned to the lessor, no adversary can distinguish the marked input of a point function from a random (non-marked) input with probability better than $1/2$, except for a negligible advantage (in the parameter λ).

We give a proof of Theorem 5 in the next section.

Proof of security

To prove the theorem, we rely on a few technical results.

Lemma 5. *Let $\alpha \in \mathbb{C}^n$ and $A_1, \ldots, A_n \in \mathbb{C}^{m \times m}$. Then, it holds that*

$$\mathrm{Tr}\Big[\sum_{i=1}^{n} \alpha_i A_i\Big] \;\leq\; \|\alpha\|_1 \cdot \sum_{i=1}^{n} |\mathrm{Tr}[A_i]|.$$

Proof. Using the Cauchy-Schwarz inequality, we have

$$\mathrm{Tr}\Big[\sum_{i=1}^{n} \alpha_i A_i\Big] = \sum_{i=1}^{n} \alpha_i \mathrm{Tr}\big[A_i\big] \leq \sqrt{\sum_{i=1}^{n} |\alpha_i|^2} \cdot \sqrt{\sum_{i=1}^{n} |\mathrm{Tr}\big[A_i\big]|^2}.$$

The claim follows from the norm inequality $\|x\|_2 \leq \|x\|_1$, for all $x \in \mathbb{C}^n$. □

Lemma 6. *Let $0 \leq \Pi \leq \mathbb{1}$ and let ϱ and σ be states such that $\mathsf{TD}(\varrho, \sigma) \leq \gamma$. Then,*

$$\mathrm{Tr}[\Pi\varrho] - \gamma \leq \mathrm{Tr}[\Pi\sigma] \leq \mathrm{Tr}[\Pi\varrho] + \gamma.$$

Proof. By the standard identity $\mathsf{TD}(\sigma, \varrho) = \max_{0 \leq \Lambda \leq \mathbb{1}} \mathrm{Tr}[\Lambda(\sigma - \varrho)]$, it follows that:

$$\begin{aligned}
\mathrm{Tr}[\Pi\sigma] &= \mathrm{Tr}[\Pi\varrho] + \mathrm{Tr}[\Pi(\sigma - \varrho)] \\
&\leq \mathrm{Tr}[\Pi\varrho] + \max_{0 \leq \Lambda \leq \mathbb{1}} \mathrm{Tr}[\Lambda(\sigma - \varrho)] \\
&= \mathrm{Tr}[\Pi\varrho] + \mathsf{TD}(\sigma, \varrho) \\
&\leq \mathrm{Tr}[\Pi\varrho] + \gamma.
\end{aligned}$$

The other inequality can be shown by reversing the role of ϱ and σ. □

Lemma 7 ([127], Lemma 18)**.** *Let $\theta \in \{0,1\}^m$ and define $\Pi_\theta^{\mathsf{eq}} = \sum_{v \in \{0,1\}^m} H^\theta \left|v\right\rangle \left\langle v\right| H^\theta \otimes H^\theta \left|v\right\rangle \left\langle v\right| H^\theta$ (i.e., the projector that checks if two registers yield the same outcome if measured in the H^θ basis). Then, the following is true for every $t \in [m]$. For any approximate EPR state,*

$$\left|\phi_{ab}^+\right\rangle = \frac{1}{\sqrt{2^m}} \sum_{v \in \{0,1\}^m} \left|v\right\rangle \otimes X^a Z^b \left|v\right\rangle,$$

where $a, b \in \{0,1\}^m$ have Hamming weight at most t, it follows that:

- $\Pi_\theta^{\mathsf{eq}} \left|\phi_{ab}^+\right\rangle = \left|\phi_{ab}^+\right\rangle$ *holds if and only if for all $i \in [m]$:*

$$(\theta_i = 0 \wedge a_i = 0) \vee (\theta_i = 1 \wedge b_i = 0).$$

- $\Pi_\theta^{\mathsf{eq}} \left|\phi_{ab}^+\right\rangle = 0$ *holds for all other cases.*

We also rely on the next lemma which is based on a result by Unruh [127, Lemma 15]. To state the lemma, we define the projector onto the subspace spanned by EPR-pairs in registers XY with up to $t \in \mathbb{N}$ single-qubit Pauli operators applied to register Y:

$$\Pi_t^{\mathsf{EPR}} = \sum_{\substack{a,b\in\{0,1\}^m \\ w(a),w(b)\leq t}} |\phi_{ab}^+\rangle\langle\phi_{ab}^+|, \qquad |\phi_{ab}^+\rangle = \frac{1}{\sqrt{2^m}} \sum_{v\in\{0,1\}^m} |v\rangle \otimes X^a Z^b |v\rangle,$$

where $w(a), w(b)$ denote the Hamming weights of the strings a and b. Since $\left\{ |\phi_{ab}^+\rangle : a, b \in \{0,1\}^m \right\}$ forms an orthogonal basis of XY, any state ϱ such that $\left(\Pi_t^{\mathsf{EPR}} \otimes \mathbb{1}_{\mathsf{R}}\right) \varrho_{\mathsf{XYR}} = \varrho_{\mathsf{XYR}}$ on registers X, Y and R can be written as follows (where a, b of weight greater than t have probability zero):

$$\varrho_{\mathsf{XYR}} = \sum_{\substack{a,b\in\{0,1\}^m \\ w(a),w(b)\leq t}} p_{ab} \left(|\phi_{ab}^+\rangle\langle\phi_{ab}^+|_{\mathsf{XY}} \otimes \sigma_{\mathsf{R}}^{a,b} \right), \tag{3.18}$$

for some arbitrary states $\sigma^{a,b}$ and indices $a, b \in \{0,1\}^m$. We show the following lemma:

Lemma 8 (Monogamy uncertainty relation)**.** *Fix a parameter $t \in \mathbb{N}$ and string $\theta \in \{0,1\}^m$. Let ϱ be a density matrix on registers X, Y and R with*

$$\left(\Pi_t^{\mathsf{EPR}} \otimes \mathbb{1}_{\mathsf{R}}\right) \varrho_{\mathsf{XYR}} \left(\Pi_t^{\mathsf{EPR}} \otimes \mathbb{1}_{\mathsf{R}}\right) = \varrho_{\mathsf{XYR}}.$$

Let $\{\Pi_{v'}\}_{v'\in\{0,1\}^m}$ be a POVM acting on register R and suppose that a measurement according to the set $\left\{H^\theta |v\rangle\langle v|_{\mathsf{X}} H^\theta \otimes \mathbb{1}_{\mathsf{Y}} \otimes \Pi_{v'\mathsf{R}}\right\}_{v'\in\{0,1\}^m}$ is performed on systems XYR. Then,

$$\Pr[v' = v] = \sum_{v\in\{0,1\}^m} \mathrm{Tr}\left[\left(H^\theta |v\rangle\langle v|_{\mathsf{X}} H^\theta \otimes \mathbb{1}_{\mathsf{Y}} \otimes \Pi_{v\mathsf{R}}\right) \varrho_{\mathsf{XYR}}\right] \leq 2^{-m}(m+1)^{2t}.$$

Hence, the min-entropy of the random variable V (with outcome v) given register R is at least

$$\mathbf{H}_{\min}(V|\mathsf{R})_\varrho \geq m - 2t\log(m+1).$$

Proof. For brevity, we define a family of projectors $\{\Lambda_u^\theta\}_u$ acting on registers X and Y, where

$$\Lambda_u^\theta = \left(H^\theta |u\rangle\langle u|_{\mathsf{X}} H^\theta \otimes \mathbb{1}_{\mathsf{Y}}\right).$$

Let T be the set of all possible indices of weight less or equal than t. Now, using decomposition (3.18), we can bound the success probability of measuring $v' = v$ using the information in the ancilla register R as follows:

$$
\begin{aligned}
&\Pr[v' = v] \\
&= \sum_{v \in \{0,1\}^m} \mathrm{Tr}\Big[\Big(H^\theta \left|v\right\rangle \left\langle v\right|_{\mathsf{X}} H^\theta \otimes \mathbb{1}_{\mathsf{Y}} \otimes \Pi_{v\mathsf{R}} \Big) \varrho_{\mathsf{XYR}} \Big] \\
&= \sum_{v \in \{0,1\}^m} \mathrm{Tr}\Big[\sum_{\substack{a,b \in \{0,1\}^m \\ w(a),w(b) \leq t}} p_{ab} \, \Big(\Lambda_v^\theta \left|\phi^+_{ab}\right\rangle \left\langle \phi^+_{ab}\right|_{\mathsf{XY}} \Lambda_v^\theta \Big) \otimes \Big(\Pi_v \sigma_{\mathsf{R}}^{a,b} \Big) \Big] && \text{(by def.)} \\
&\leq \sum_{v \in \{0,1\}^m} \Big(\sum_{\substack{a,b \in \{0,1\}^m \\ w(a),w(b) \leq t}} p_{ab} \Big) \cdot \Big(\sum_{\substack{a,b \in \{0,1\}^m \\ w(a),w(b) \leq t}} \| \Lambda_v^\theta \left|\phi^+_{ab}\right\rangle \|^2 \cdot \mathrm{Tr}\big[\Pi_v \sigma_{\mathsf{R}}^{a,b} \big] \Big) && \text{(Lem. 5)} \\
&= \sum_{v \in \{0,1\}^m} \sum_{\substack{a,b \in \{0,1\}^m \\ w(a),w(b) \leq t}} \| H^\theta \left|v\right\rangle \left\langle v\right|_{\mathsf{X}} H^\theta \otimes \mathbb{1}_{\mathsf{Y}} \left|\phi^+_{ab}\right\rangle \|^2 \cdot \mathrm{Tr}\big[\Pi_v \sigma_{\mathsf{R}}^{a,b} \big] && \text{(by def.)} \\
&= \sum_{v \in \{0,1\}^m} \sum_{\substack{a,b \in \{0,1\}^m \\ w(a),w(b) \leq t}} \| H^\theta \left|v\right\rangle \left\langle v\right|_{\mathsf{X}} H^\theta \otimes X^a Z^b_{\mathsf{Y}} \left|\phi^+\right\rangle \|^2 \cdot \mathrm{Tr}\big[\Pi_v \sigma_{\mathsf{R}}^{a,b} \big] \\
&= \sum_{v \in \{0,1\}^m} \sum_{\substack{a,b \in \{0,1\}^m \\ w(a),w(b) \leq t}} \| H^\theta \otimes X^a Z^b H^\theta \big(\left|v\right\rangle \left\langle v\right|_{\mathsf{X}} \otimes \mathbb{1}_{\mathsf{Y}} \big) \left|\phi^+\right\rangle \|^2 \cdot \mathrm{Tr}\big[\Pi_v \sigma_{\mathsf{R}}^{a,b} \big] \\
&= \sum_{v \in \{0,1\}^m} \sum_{\substack{a,b \in \{0,1\}^m \\ w(a),w(b) \leq t}} \frac{\mathrm{Tr}\big[\Pi_v \sigma_{\mathsf{R}}^{a,b} \big]}{2^m} = \sum_{\substack{a,b \in \{0,1\}^m \\ w(a),w(b) \leq t}} \frac{\mathrm{Tr}\big[\sigma_{\mathsf{R}}^{a,b} \big]}{2^m} = \frac{|T|}{2^m},
\end{aligned}
$$

where in the second-to-last step we used the completeness property that $\sum_v \Pi_v = \mathbb{1}$, and in the last step we use that the $\sigma^{a,b}$ have unit trace, for every $a, b \in \{0,1\}^m$. It now suffices to bound $|T|$, the number of error indices of weight less or equal to t. In total we have t indices to assign to $m+1$ possible choices (we add an additional degree of freedom to account for when there are no errors assigned). Since we have two independent indices $a, b \in \{0,1\}^m$, we get:

$$\Pr[v' = v] \leq 2^{-m}|T| \leq 2^{-m}(m+1)^{2t}.$$

This proves the claim. □

Let us now proceed with the security proof. We consider the following sequence of hybrids of $\mathsf{SSLGame}$. We will show that the optimal winning probability in each successive hybrid changes at most negligibly. We will then bound the optimal winning probability in the final hybrid.

H_0: This is the original game $\mathsf{SSLGame}$ in Section 3.5:

- The lessor runs $\mathsf{SSL.Lease}(1^\lambda, y \in \{0,1\}^\lambda)$ to sample $v \leftarrow \{0,1\}^m$ and $\theta \leftarrow G(y) \in \{0,1\}^m$, and sends $(\left|v^\theta\right\rangle, H(v))$ together with a circuit for $\mathsf{SSL.Eval}$ to the adversary $\mathcal{A}$.

- Having access to the random oracles G and H, the adversary $\mathcal{A}$ outputs a (possibly entangled) state σ on two registers Y and R, and sends the register Y to the lessor.
- For verification, the lessor runs $\mathsf{SSL.Verify}(y, \mathsf{Y})$: Measure the register Y in the H^{θ} basis according to $\theta = G(y)$. If the outcome is equal to v such that $H(v) = z$, the lessor outputs $\mathsf{ok} = 1$ and lets the game continue, otherwise, the lessor outputs $\mathsf{ok} = 0$ and $\mathcal{A}$ loses.
- Conditioning on $\mathsf{ok} = 1$, the lessor sends the adversary a sample $x \leftarrow \mathcal{D}_y$ to which $\mathcal{A}$ responds with a bit (we refer to this phase of the security game as the "input challenge phase"). Using the string y given as input, the lessor outputs 1, if the bit is equal to $P_y(x)$, and 0 otherwise.

H_1: The game is the same as before, except that in the input challenge phase the lessor samples $x \leftarrow D_y$, and sends $G(x)$ to $\mathcal{A}$, (instead of sending x directly).

H_2: The game is the same as before, except for the input challenge phase. The lessor samples $x \leftarrow D_y$. Then, if $x \neq y$, the lessor chooses $\theta' \leftarrow \{0,1\}^m$ and sends θ' to $\mathcal{A}$ (instead of $G(x)$).

H_3: The game is the same as before, except that the lessor samples $\theta \leftarrow \{0,1\}^m$ (instead of $\theta \leftarrow G(y)$). Then, in the input challenge phase, the lessor samples $x \leftarrow D_y$. If $x = y$, the lessor sends θ to $\mathcal{A}$.

H_4: The game is identical to the game before, except that we replace $H(v)$ with a uniformly random string $z \leftarrow \{0,1\}^{\lambda}$.

First, we show that the advantage of any adversary in H_4 is negligible. In the rest of the section, we denote by $p(H_i)$ the optimal winning probability in hybrid H_i.

Lemma 9. $p(H_4) \leq \frac{1}{2}$.

Proof. First, the optimal probability of the adversary winning the game can only increase if we remove the verification portion of the game, and the lessor directly executes the input challenge phase.

Then, we consider the state received by the adversary in the two distinct cases of the input challenge phase.

- The lessor samples the marked point. In this case, the state received by the adversary is the following, which is completely independent of the oracle H:

  $$\mathbb{E}_{\theta,v}\left(|v^{\theta}\rangle\langle v^{\theta}|\otimes|\theta\rangle\langle\theta|\right)\otimes\mathbb{E}_z\,|z\rangle\langle z|\ .$$

  Notice that the latter state is maximally mixed.

- The lessor samples a point other than the marked point. In this case, the adversary receives the following state, which is again independent of the oracle:

  $$\mathbb{E}_{\theta,\theta',v}\left(|v^{\theta}\rangle\langle v^{\theta}|\otimes|\theta'\rangle\langle\theta'|\right)\otimes\mathbb{E}_z\,|z\rangle\langle z|\ .$$

  The latter state is again maximally mixed.

Thus, an adversary can win the game H_4 with probability at most $\frac{1}{2}$. □

We will now show that the optimal success probabilities in successive hybrids do not deviate by more than a negligible amount.

Lemma 10. $|p(H_1)-p(H_0)|=\mathsf{negl}(\lambda)$.

Proof. The proof follows immediately from the fact that G is a random oracle, and hence the pre-image x does not help the adversary and can be simulated. □

Lemma 11. $|p(H_2)-p(H_1)|=\mathsf{negl}(\lambda)$.

Proof. This follows immediately from the following observation: Any adversary that wins with non-negligible difference in H_2 and H_1 immediately yields a distinguisher for $G(X_\lambda)$ and $U_{m(\lambda)}$. This violates Corollary 1. □

Lemma 12. $|p(H_3)-p(H_2)|=\mathsf{negl}(\lambda)$.

Proof. The proof is analogous as in the previous lemma, where an adversary that wins with probabilities that differ non-negligibly in H_3 and H_2 yields a distinguisher for $G(X_\lambda)$ and $U_{m(\lambda)}$. □

The crux of the security proof is showing that $p(H_3)$ and $p(H_4)$ are negligibly close.

Lemma 13. $|p(H_4)-p(H_3)|=\mathsf{negl}(\lambda)$.

The rest of the section is devoted to proving this lemma. At a high level, the proof has two parts:

- For any adversary making q queries to the oracle, we bound the difference between the winning probability in H_3 and in H_4 by $\mathsf{poly}(q) \cdot M$, where M is a quantity related to the probability that the adversary queries the oracle at the encoded string v.
- Then, we show that the quantity M is negligible.

Lemma 14. *Let $\mathcal{A}$ be an adversary for H_3 and H_4, making $\mathsf{poly}(\lambda)$ oracle queries (pre and post verification). Suppose that $\mathcal{A}$ passes the verification step with probability at least $\frac{1}{2} - \mathsf{negl}(\lambda)$ in H_3. Let $\mathcal{A}$ be specified by the unitary U (i.e. $\mathcal{A}$ alternates oracles calls with applications of U). Let $p_{v,\theta,z,H} \in [0,1]$, and let $\varrho_{\mathsf{R}}^{v,\theta,z,H}$ be density matrices, for all v, θ, z, H. Let*

$$\sigma_{\mathsf{LR}} = \mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \left(|H\rangle \langle H| \otimes |v\rangle \langle v| \otimes |\theta\rangle \langle\theta| \otimes |z\rangle \langle z| \right)_{\mathsf{L}} \otimes \varrho_{\mathsf{R}}^{v,\theta,z,H}$$

be the post-verification state of the lessor and $\mathcal{A}$ in H_4 conditioned on $\mathcal{A}$ passing the verification step. Let $\tau_\theta = \frac{1}{2} |\theta\rangle \langle\theta| + \frac{1}{2}\mathbb{E}_{\theta'} |\theta'\rangle \langle\theta'|$. Then,

$$|\Pr[\mathcal{A} \text{ wins in } H_3] - \Pr[\mathcal{A} \text{ wins in } H_4]| \leq \mathsf{poly}(\lambda) \cdot M + \mathsf{negl}(\lambda)\,,$$

where

$$\begin{aligned} M = &\frac{1}{2}\, \mathbb{E}_H \mathbb{E}_v \mathbb{E}_\theta \mathbb{E}_z \mathbb{E}_k p_{v,\theta,z,H} \mathrm{Tr} |v\rangle \langle v| \, (UO^{H_{v,z}})^k \left(\varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes \tau_\theta \right) \left((UO^{H_{v,z}})^k \right)^\dagger \\ &+ \frac{1}{2}\, \mathbb{E}_H \mathbb{E}_v \mathbb{E}_\theta \mathbb{E}_z \mathbb{E}_k p_{v,\theta,z,H} \mathrm{Tr} |v\rangle \langle v| \, (UO^{H})^k \left(\varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes \tau_\theta \right) \left((UO^{H})^k \right)^\dagger . \end{aligned}$$

Proof. As we have done in several earlier proofs, we can recast H_3 as follows: $\mathcal{A}$ receives a uniformly random z, and gets access to a the reprogrammed oracle $H_{v,z}$. Let $|v^\theta\rangle$ denote the encoding of string v using basis θ. Let q_1 and q_2 denote the number of queries made by the adversary, respectively, before and after the verification phase.

First notice that the global states of the lessor and adversary right before the verification is executed are negligibly close in trace distance in H_3 and H_4.

$$\begin{aligned} &\mathbb{E}_H \mathbb{E}_v \mathbb{E}_\theta \mathbb{E}_z \, |H\rangle \langle H| \otimes |v\rangle \langle v| \otimes |\theta\rangle \langle \theta| \otimes \left((UO^{H_{v,z}})^{q_1} \, |v^\theta\rangle \langle v^\theta| \otimes |z\rangle \langle z| \left((UO^{H_{v,z'}})^{q_1} \right)^\dagger \right) \\ \approx &\mathbb{E}_H \mathbb{E}_v \mathbb{E}_\theta \mathbb{E}_z \, |H\rangle \langle H| \otimes |v\rangle \langle v| \otimes |\theta\rangle \langle \theta| \otimes \left((UO^{H})^{q_1} \, |v^\theta\rangle \langle v^\theta| \otimes |z\rangle \langle z| \left((UO^{H})^{q_1} \right)^\dagger \right) . \end{aligned} \tag{3.19}$$

Here we have stored the complete function H in an additional register, the quantum way of formulating indistinguishability of the joint distribution of H and the adversary's state.

Equation (3.19) follows from the one-way-to-hiding lemma (Lemma 4), and the fact that $\mathcal{A}$ only queries at v with negligible probability (otherwise $\mathcal{A}$ would straightforwardly imply an adversary that wins the monogamy game (more precisely the variant of Lemma 1).

It follows that:

- The probabilities of $\mathcal{A}$ passing the verification step in H_3 and in H_4 are negligibly close.
- The post-verification states, conditioned on passing verification must be negligibly close (this uses (3.19) together with the fact that, by hypothesis, $\mathcal{A}$ passes verification with probability at least $\frac{1}{2} - \mathsf{negl}(\lambda)$).

By definition, the joint state of lessor and adversary post-verification state in H_4 conditioned on $\mathcal{A}$ passing verification is

$$\sigma_{\mathsf{LR}} = \mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \left(|H\rangle \langle H| \otimes |v\rangle \langle v| \otimes |\theta\rangle \langle \theta| \otimes |z\rangle \langle z| \right)_{\mathsf{L}} \otimes \varrho_{\mathsf{R}}^{v,\theta,z,H} .$$

Let the analogous state in H_3 be

$$\tilde{\sigma}_{\mathsf{LR}} = \mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \left(|H\rangle \langle H| \otimes |v\rangle \langle v| \otimes |\theta\rangle \langle \theta| \otimes |z\rangle \langle z| \right)_{\mathsf{L}} \otimes \tilde{\varrho}_{\mathsf{R}}^{v,\theta,z,H} .$$

Then $\sigma_{\mathsf{L,R}} \approx \tilde{\sigma}_{\mathsf{L,R}}$. Now, denote by $\{\Pi^0, \Pi^1\}$ the projective measurement performed by $\mathcal{A}$ to guess the answer to the input challenge phase. Then,

$$\begin{aligned}
&\Pr[\mathcal{A} \text{ wins in } H_4 | \text{verification is passed}] \\
&= \mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \Bigg[\frac{1}{2} \mathrm{Tr} \Pi^1 (UO^H)^{q_2} \varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta\rangle \langle \theta| \left((UO^H)^{q_2} \right)^\dagger \\
&+ \frac{1}{2} \mathbb{E}_{\theta'} \mathrm{Tr} \Pi^0 (UO^H)^{q_2} \varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta'\rangle \langle \theta'| \left((UO^H)^{q_2} \right)^\dagger \Bigg] .
\end{aligned} \tag{3.20}$$

And, similarly,

$$\begin{aligned}
&\Pr[\mathcal{A} \text{ wins in } H_3 | \text{verification is passed}] \\
&= \mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \Bigg[\frac{1}{2} \mathrm{Tr} \Pi^1 (UO^{H_{v,z}})^{q_2} \tilde{\varrho}_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta\rangle \langle \theta| \left((UO^{H_{v,z}})^{q_2} \right)^\dagger \\
&+ \frac{1}{2} \mathbb{E}_{\theta'} \mathrm{Tr} \Pi^0 (UO^{H_{v,z}})^{q_2} \tilde{\varrho}_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta'\rangle \langle \theta'| \left((UO^{H_{v,z}})^{q_2} \right)^\dagger \Bigg] \\
&\approx \mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \Bigg[\frac{1}{2} \mathrm{Tr} \Pi^1 (UO^{H_{v,z}})^{q_2} \varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta\rangle \langle \theta| \left((UO^{H_{v,z}})^{q_2} \right)^\dagger \\
&+ \frac{1}{2} \mathbb{E}_{\theta'} \mathrm{Tr} \Pi^0 (UO^{H_{v,z}})^{q_2} \varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta'\rangle \langle \theta'| \left((UO^{H_{v,z}})^{q_2} \right)^\dagger \Bigg] .
\end{aligned} \tag{3.21}$$

Using equations (3.20) and (3.21), and applying the O2H lemma twice (once to bound the distance between the first terms in expressions (3.20) and (3.21), and once to bound the distance between the second terms in (3.20) and (3.21)), we obtain:

$$\begin{aligned}
&|\Pr[\mathcal{A} \text{ wins in } H_4|\text{verification is passed}] - \Pr[\mathcal{A} \text{ wins in } H_3|\text{verification is passed}]| \\
&\leq \mathsf{poly}(\lambda) \cdot \frac{1}{2}\mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \mathrm{Tr}|v\rangle\langle v|\, (UO^{H})^k \left(\varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta\rangle\langle\theta|\right) \left((UO^{H})^k\right)^{\dagger} \\
&+ \mathsf{poly}(\lambda) \cdot \frac{1}{2}\mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \mathrm{Tr}|v\rangle\langle v|\, (UO^{H_{v,z}})^k \left(\varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta\rangle\langle\theta|\right) \left((UO^{H_{v,z}})^k\right)^{\dagger} \\
&+ \mathsf{poly}(\lambda) \cdot \frac{1}{2}\mathbb{E}_{v,\theta,z,H,\theta'}\, p_{v,\theta,z,H} \mathrm{Tr}|v\rangle\langle v|\, (UO^{H})^k \left(\varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta'\rangle\langle\theta'|\right) \left((UO^{H})^k\right)^{\dagger} \\
&+ \mathsf{poly}(\lambda) \cdot \frac{1}{2}\mathbb{E}_{v,\theta,z,H,\theta'}\, p_{v,\theta,z,H} \mathrm{Tr}|v\rangle\langle v|\, (UO^{H_{v,z}})^k \left(\varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes |\theta'\rangle\langle\theta'|\right) \left((UO^{H_{v,z}})^k\right)^{\dagger} + \mathsf{negl}(\lambda) \\
&= \mathsf{poly}(\lambda) \cdot \frac{1}{2}\mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \mathrm{Tr}|v\rangle\langle v|\, (UO^{H})^k \left(\varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes \tau_\theta\right) \left((UO^{H})^k\right)^{\dagger} \\
&+ \mathsf{poly}(\lambda) \cdot \frac{1}{2}\mathbb{E}_{v,\theta,z,H}\, p_{v,\theta,z,H} \mathrm{Tr}|v\rangle\langle v|\, (UO^{H_{v,z}})^k \left(\varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes \tau_\theta\right) \left((UO^{H_{v,z}})^k\right)^{\dagger} + \mathsf{negl}(\lambda) \\
&= \mathsf{poly}(\lambda) \cdot M + \mathsf{negl}(\lambda)\,, \qquad (3.22)
\end{aligned}$$

where to get two equalities we used the definition of τ_θ and M. This is the desired bound. □

In the rest of the section, we show that the quantity M from Lemma 14 is negligible. First of all, notice that M is negligible if and only if the second term in M is negligible, i.e., if and only if,

$$\mathbb{E}_H\mathbb{E}_v\mathbb{E}_\theta\mathbb{E}_z\mathbb{E}_k p_{v,\theta} \mathrm{Tr}|v\rangle\langle v|\, (UO^{H})^k \left(\varrho_{\mathsf{R}}^{v,\theta,z,H} \otimes \tau_\theta\right) \left((UO^{H})^k\right)^{\dagger} = \mathsf{negl}(\lambda)\,. \qquad (3.23)$$

where we are using the same notation as in Lemma 14. Thus, what we wish to show is equivalent to showing that, for any adversary $\mathcal{A}$ in H_4 who passes verification with probability at least $\frac{1}{2} - \mathsf{negl}(\lambda)$, the probability of querying the oracle at the encoded string v at any point after a successful verification is negligible.

Thus, we will show that (3.23) is negligible. First, notice that an adversary $\mathcal{A}$ which passes verification in H_4 with probability at least $\frac{1}{2} - \mathsf{negl}(\lambda)$ and violates (3.23) immediately implies an adversary which succeeds at the following game $\widetilde{H}_0$ with non-negligible probability.

$\widetilde{H}_0$: This is identical to H_4 except we ask the adversary to return a guess

v' for the encoded string v, instead of a bit. $\mathcal{A}$ wins if $v' = v$.

The reduction crucially uses the hypothesis that $\mathcal{A}$ passes verification with probability at least $\frac{1}{2} - \mathsf{negl}(\lambda)$. We will show through another sequence of hybrids (which we denote using tildes) that the optimal winning probability in $\widetilde{H}_0$ is negligible. This will complete the proof that the quantity in (3.23), and thus M is negligible, for any adversary $\mathscr{A}$ who passes verification with probability at least $\frac{1}{2} - \mathsf{negl}(\lambda)$. Since the optimal winning probability in H_3 and H_4 is at least $\frac{1}{2}$ (the honest strategy followed by random guessing achieves $\frac{1}{2}$), this concludes the proof of Lemma 13, and hence that the optimal winning probability in H_0 is at most $\frac{1}{2} + \mathsf{negl}(\lambda)$. The following are the hybrids.

$\widetilde{H}_1$: Instead of sampling $v \leftarrow \{0,1\}^m$ and $\theta \leftarrow \{0,1\}^m$ at the beginning of the game, the lessor now prepares an EPR pair on two registers X and Y, and sends the registers YZ of the state $|\phi^+\rangle_{\mathsf{XY}} \otimes |z\rangle_{\mathsf{Z}}$ to $\mathscr{A}$. Rather than running SSL.Verify for verification and measuring the register Y, the lessor now measures both registers X and Y in the H^θ basis for a random $\theta \leftarrow \{0,1\}^m$, and checks if the outcomes result in the same string, which we denote by v.

$\widetilde{H}_2$: This game is identical to the one before, except that we change the verification procedure as follows. Instead of measuring each of the registers X and Y in the H^θ basis, the lessor now measures a bipartite projector Π_θ^{eq} in order to check if the registers XY yield the same outcome if measured in the H^θ basis. We define the projector as follows:

$$\Pi_\theta^{\mathsf{eq}} = \sum_{v \in \{0,1\}^m} H^\theta |v\rangle \langle v|_{\mathsf{X}} H^\theta \otimes H^\theta |v\rangle \langle v|_{\mathsf{Y}} H^\theta.$$

Afterwards, the lessor measures register X in the H^θ basis to determine v.

We will denote these hybrids using a tilde to distinguish them from the original sequence of hybrids.

Lemma 15. $p(\widetilde{H}_1) = p(\widetilde{H}_0)$.

Proof. The argument is fairly standard. We consider the following two statements:

- sample $v \leftarrow \{0,1\}^m$, let $\theta \in \{0,1\}^m$, and output $\bigotimes_{i=1}^m |v_i^{\theta_i}\rangle_{\mathsf{Y}}$.
- create an m-qubit EPR pair $|\phi^+\rangle_{\mathsf{XY}}$, measure X in the H^θ basis, and output register Y.

It is evident that the equivalence of the two statements implies that $p(\widetilde{H}_1)$ and $p(\widetilde{H}_0)$ are identical. Note that we omit the register $|z\rangle$ in the proof, since it is independent of the EPR registers and thus does not affect the argument. Consider the following family of projectors given by

$$\{(H^\theta |v\rangle \langle v| H^\theta \otimes \mathbb{1}_{\mathsf{Y}})\}_{v \in \{0,1\}^m}.$$

Let us now analyze the post-measurement state $|\psi_v\rangle / \sqrt{\langle \psi_v \,|\, \psi_v\rangle}$ with respect to the state given by $|\psi_v\rangle = \left(H^\theta |v\rangle \langle v|_{\mathsf{X}} H^\theta \otimes \mathbb{1}_{\mathsf{Y}}\right) |\phi^+\rangle$. We have,

$$\begin{aligned}
|\psi_v\rangle_{\mathsf{XY}} &= \left(H^\theta |v\rangle \langle v| H^\theta \otimes \mathbb{1}\right) |\phi^+\rangle_{\mathsf{XY}} \\
&= \Big(\big(H^\theta \otimes \mathbb{1}\big)\big(|v\rangle \langle v| \otimes \mathbb{1}\big)\big(H^\theta \otimes \mathbb{1}\big)\Big) |\phi^+\rangle_{\mathsf{XY}} \\
&= \Big(\big(H^\theta \otimes \mathbb{1}\big)\big(|v\rangle \langle v| \otimes \mathbb{1}\big)\big(\mathbb{1} \otimes H^\theta\big)\Big) |\phi^+\rangle_{\mathsf{XY}} \qquad \text{(ricochet property)} \\
&= 2^{-m/2} \sum_{v' \in \{0,1\}^m} \Big(\big(H^\theta \otimes \mathbb{1}\big)\big(|v\rangle \langle v| \otimes \mathbb{1}\big)\big(\mathbb{1} \otimes H^\theta\big)\Big) |v'\rangle_{\mathsf{X}} \otimes |v'\rangle_{\mathsf{Y}} \\
&= 2^{-m/2} \sum_{v' \in \{0,1\}^m} H^\theta |v\rangle_{\mathsf{X}} \langle v \,|\, v'\rangle \otimes H^\theta |v'\rangle_{\mathsf{Y}} \\
&= 2^{-m/2} H^\theta |v\rangle_{\mathsf{X}} \otimes H^\theta |v\rangle_{\mathsf{Y}} \,.
\end{aligned}$$

This proves the claim, since the Y register of $|\psi_v\rangle / \sqrt{\langle \psi_v \,|\, \psi_v\rangle}$ is identical to $\bigotimes_{i=1}^m |v_i^{\theta_i}\rangle$. □

Lemma 16. $p(\widetilde{H}_2) = p(\widetilde{H}_1)$

Proof. The lemma is immediate as the measurement in $\tilde{H}_2$ is a coarse-graining of the measurement in $\tilde{H}_1$, with the acceptance condition remaining the same. □

In the remaining part of the proof, we will show that $p(\widetilde{H}_2)$ is negligible. The following is an important technical lemma, which is inspired by Lemma 16 and Lemma 19 in [127].

Lemma 17. $p(\widetilde{H}_2) = \mathsf{negl}(\lambda)$.

Proof. Let $\mathcal{A}$ be an adversary for $\widetilde{H}_2$. Denote by v' the final guess returned by the adversary, and by v the encoded string. Let ok be a random variable for whether the verification passes. Then, the winning probability of $\mathcal{A}$ in $\widetilde{H}_2$ is given by:

$$\Pr\left[v' = v \;\wedge\; \mathsf{ok} = 1\right].$$

We show that, for any $t \in [m]$,

$$\Pr\left[v' = v \;\wedge\; \mathsf{ok} = 1\right] \;\le\; 2^{-m}(m+1)^{2t} + 2^{\frac{-t-1}{2}}. \tag{3.24}$$

Picking $t \approx \sqrt{m}$ then gives the desired result, as the RHS becomes negligible in λ.

Fix a basis choice $\theta \in \{0,1\}^m$. Let ϱ_θ be the state on registers X, Y and R in $\widetilde{H}_2$ after the verification, where R is the leftover register held onto by $\mathcal{A}$ that also includes the challenge τ_θ (where τ_θ was defined in Lemma 14) sent by the lessor after verification.

In the analysis that follows, it is convenient to approximate ϱ_θ by an ideal state that is diagonal in a basis for the image of $\Pi_t^{\mathsf{EPR}} \otimes \mathbb{1}_{\mathsf{R}}$, where Π_t^{EPR} is as defined in Lemma 8. Recall that Π_t^{EPR} projects onto the subspace spanned by EPR pairs with up to t Pauli errors, i.e., onto the space spanned by the orthogonal basis states $\left\{ |\phi_{ab}^+\rangle : a, b \in \{0,1\}^m \right\}$, where

$$|\phi_{ab}^+\rangle = \frac{1}{\sqrt{2^m}} \sum_{v \in \{0,1\}^m} |v\rangle \otimes X^a Z^b |v\rangle . \tag{3.25}$$

We can use Lemma 3 to argue that there exists such an ideal state $\varrho_\theta^{\mathsf{id}}$, and that the trace distance between the two states satisfies:

$$\|\varrho_\theta - \varrho_\theta^{\mathsf{id}}\|_{\mathrm{tr}} \leq \sqrt{1 - \mathrm{Tr}\left[\left(\Pi_t^{\mathsf{EPR}} \otimes \mathbb{1}_{\mathsf{R}}\right) \varrho_\theta\right]} .$$

We can represent the adversary's strategy in guessing v, after verification, by a projective measurement $\{\Pi_{v'}\}_{v'}$.

We are now ready to bound the probability $\Pr\left[v' = v \ \wedge \ \mathsf{ok} = 1\right]$. Let Θ be a random variable for the basis choice made by the lessor. Then, by marginalizing over Θ, we get:

$$\begin{aligned}
\Pr\left[v' = v \ \wedge \ \mathsf{ok} = 1\right] &= \sum_{\theta \in \{0,1\}^m} 2^{-m} \cdot \Pr\left[v' = v | \, \mathsf{ok} = 1 \ \wedge \ \Theta = \theta\right] \cdot \Pr[\mathsf{ok} = 1 | \Theta = \theta] \\
&\leq \sum_{\theta \in \{0,1\}^m} 2^{-m} \cdot \Pr\left[v' = v | \, \mathsf{ok} = 1 \ \wedge \ \Theta = \theta\right] \\
&= \mathbb{E}_\theta \Pr[v' = v \,|\, \mathsf{ok} = 1 \ \wedge \ \Theta = \theta].
\end{aligned} \tag{3.26}$$

Fix any θ. Using Lemma 6 and Lemma 8 we obtain:

$$\begin{aligned}
\Pr\left[v' = v \,|\, \mathsf{ok} = 1 \wedge \Theta = \theta\right] &\leq 2^{-m}(m+1)^{2t} + \mathsf{TD}(\varrho_\theta, \varrho_\theta^{\mathsf{id}}) \\
&\leq 2^{-m}(m+1)^{2t} + \sqrt{1 - \mathrm{Tr}\left[\left(\Pi_t^{\mathsf{EPR}} \otimes \mathbb{1}_{\mathsf{R}}\right) \varrho_\theta\right]} .
\end{aligned} \tag{3.27}$$

Now, averaging over θ in the above inequality gives:

$$\begin{aligned}
\mathbb{E}_\theta \Pr[v' = v \,|\, \mathsf{ok} = 1 \wedge \Theta = \theta] &\leq 2^{-m}(m+1)^{2t} + \mathbb{E}_\theta \sqrt{1 - \mathrm{Tr}\left[\left(\Pi_t^{\mathsf{EPR}} \otimes \mathbb{1}_{\mathsf{R}}\right) \varrho_\theta\right]} \\
&\leq 2^{-m}(m+1)^{2t} + \sqrt{\mathbb{E}_\theta \mathrm{Tr}\left[\left(\left(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\right) \otimes \mathbb{1}_{\mathsf{R}}\right) \varrho_\theta\right]} .
\end{aligned} \tag{3.28}$$

where the last inequality follows from Jensen's inequality. We will proceed to bound the above term $\mathbb{E}_\theta \mathrm{Tr}\left[\left(\left(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\right) \otimes \mathbb{1}_{\mathsf{R}}\right) \varrho_\theta\right]$ by 2^{-t-1}. Let us first show that for any $a, b \in \{0,1\}^m$:

$$p_{ab} \overset{\mathrm{def}}{=} \sum_{\theta \in \{0,1\}^m} 2^{-m} \mathrm{Tr}\left[\left(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\right) \Pi_\theta^{\mathsf{eq}} \, |\phi_{ab}^+\rangle \langle \phi_{ab}^+|_{\mathsf{XY}}\right] \ \leq \ 2^{-t-1}. \tag{3.29}$$

This follows from considering the following two cases:

- $w(a), w(b) \leq t$: Using Lemma 7 we find that one of the following is true. Depending on θ, either $\Pi_\theta^{\mathsf{eq}} |\phi_{ab}^+\rangle = 0$ or $\Pi_\theta^{\mathsf{eq}} |\phi_{ab}^+\rangle = |\phi_{ab}^+\rangle$. We also get that $\big(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\big) |\phi_{ab}^+\rangle = 0$, since $\Pi_t^{\mathsf{EPR}} |\phi_{ab}^+\rangle = |\phi_{ab}^+\rangle$, and thus it follows that $p_{ab} = 0$.
- $\max\big(w(a), w(b)\big) \geq t+1$: Here, Lemma 7 implies that there are at most $2^m/2^{t+1}$ many values of θ for which it holds that $\Pi_\theta^{\mathsf{eq}} |\phi_{ab}^+\rangle \neq 0$, and thus $p_{ab} \leq 2^{-m} \cdot 2^m/2^{t+1} = 2^{-t-1}$.

Observe now that Π_t^{EPR} and $|\phi_{ab}^+\rangle\langle\phi_{ab}^+|$ are diagonal in the Bell basis, hence they commute. Lemma 7 implies that the same is also true for the projector Π_θ^{eq}. For any fixed $\theta \in \{0,1\}^m$, we express ϱ_θ as a generic density operator on registers X, Y and R such that, for a finite index set I^θ, coefficients q_{ij} and an orthogonal basis $\{|\psi^{i,\theta}\rangle \,:\, i \in I^\theta\}$ the registers X and Y:

$$\varrho_\theta = \sum_{i,j \in I^\theta} q_{ij}\ |\psi^{j,\theta}\rangle\langle\psi^{j,\theta}|_{\mathsf{XY}} \otimes \sigma_{\mathsf{R}}^{i,j,\theta}, \tag{3.30}$$

where $\sigma^{i,j,\theta}$ are matrices for indices $i, j \in I^\theta$. Since we assumed that ϱ_θ is the state conditioned on the verification being successful for some θ, we have the property that

$$\big(\Pi_\theta^{\mathsf{eq}} \otimes \mathbb{1}_{\mathsf{R}}\big)\varrho_\theta\big(\Pi_\theta^{\mathsf{eq}} \otimes \mathbb{1}_{\mathsf{R}}\big) = \varrho_\theta, \qquad \forall \theta \in \{0,1\}^m. \tag{3.31}$$

In other words, ϱ_θ on is invariant under the action of the projector $\Pi_\theta^{\mathsf{eq}} \otimes \mathbb{1}_{\mathsf{R}}$. Then,

$$\begin{aligned}
&\mathbb{E}_\theta \mathrm{Tr}\Big[\big(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\big) \otimes \mathbb{1}_{\mathsf{R}}\, \varrho_\theta\Big] \\
&= \sum_{\theta \in \{0,1\}^m} 2^{-m} \mathrm{Tr}\big[\big(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\big) \otimes \mathbb{1}_{\mathsf{R}}\, \varrho_\theta\big] \\
&= \sum_{\theta \in \{0,1\}^m} 2^{-m} \mathrm{Tr}\big[\big(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\big) \otimes \mathbb{1}_{\mathsf{R}}\, \big(\Pi_\theta^{\mathsf{eq}} \otimes \mathbb{1}_{\mathsf{R}}\big)\varrho_\theta\big(\Pi_\theta^{\mathsf{eq}} \otimes \mathbb{1}_{\mathsf{R}}\big)\big] && \text{(Eq. (3.31))} \\
&= \sum_{\theta \in \{0,1\}^m} 2^{-m} \mathrm{Tr}\Big[\Big(\big(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\big)\Pi_\theta^{\mathsf{eq}} \otimes \mathbb{1}_{\mathsf{R}}\Big)\, \varrho_\theta\Big] \\
&= \sum_{\theta \in \{0,1\}^m} 2^{-m} \mathrm{Tr}\Big[\Big(\sum_{a,b \in \{0,1\}^m} |\phi_{ab}^+\rangle\langle\phi_{ab}^+|\Big)\big(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\big)\Pi_\theta^{\mathsf{eq}} \otimes \mathbb{1}_{\mathsf{R}}\, \varrho_\theta\Big] \\
&= \sum_{\theta \in \{0,1\}^m} \sum_{a,b \in \{0,1\}^m} 2^{-m}\, \mathrm{Tr}\Big[\,|\phi_{ab}^+\rangle\langle\phi_{ab}^+|\, \big(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\big)\Pi_\theta^{\mathsf{eq}} \otimes \mathbb{1}_{\mathsf{R}}\, \varrho_\theta\Big] \\
&= \sum_{\theta \in \{0,1\}^m} \sum_{a,b \in \{0,1\}^m} 2^{-m}\, \mathrm{Tr}\Big[\big(\mathbb{1} - \Pi_t^{\mathsf{EPR}}\big)\Pi_\theta^{\mathsf{eq}} \otimes \mathbb{1}_{\mathsf{R}}\, \big(\,|\phi_{ab}^+\rangle\langle\phi_{ab}^+| \otimes \mathbb{1}_{\mathsf{R}}\big)\varrho_\theta\big(\,|\phi_{ab}^+\rangle\langle\phi_{ab}^+| \otimes \mathbb{1}_{\mathsf{R}}\big)\Big].
\end{aligned}$$

In the third to last line, we inserted the complete set $\sum_{a,b} |\phi_{ab}^+\rangle\langle\phi_{ab}^+| = \mathbb{1}$. Then, using the definition of ϱ in Eq.(3.30), we can continue to expand the expression above as follows:

$$\begin{aligned}
&\sum_{\theta\in\{0,1\}^m}\sum_{a,b\in\{0,1\}^m} 2^{-m}\,\mathrm{Tr}\Big[\big(\mathbb{1}-\Pi_t^{\mathsf{EPR}}\big)\Pi_\theta^{\mathsf{eq}}\otimes\mathbb{1}_{\mathsf{R}}\,\big(\,|\phi^+_{ab}\rangle\langle\phi^+_{ab}|\otimes\mathbb{1}_{\mathsf{R}}\big)\varrho_\theta\big(\,|\phi^+_{ab}\rangle\langle\phi^+_{ab}|\otimes\mathbb{1}_{\mathsf{R}}\big)\Big]\\
&=\sum_{\theta\in\{0,1\}^m}\sum_{a,b\in\{0,1\}^m} 2^{-m}\sum_{i,j\in I^\theta} q_{ij}\,\mathrm{Tr}\Big[\big(\mathbb{1}-\Pi_t^{\mathsf{EPR}}\big)\Pi_\theta^{\mathsf{eq}}\;|\phi^+_{ab}\rangle\,\langle\phi^+_{ab}|\psi^{i,\theta}\rangle\,\langle\psi^{j,\theta}|\phi^+_{ab}\rangle\,\langle\phi^+_{ab}|\otimes\sigma_{\mathsf{R}}^{i,j,\theta}\Big]\\
&=\sum_{a,b\in\{0,1\}^m}\sum_{i,j\in I^\theta} p_{ab}\,q_{ij}\,\langle\phi^+_{ab}|\psi^{i,\theta}\rangle\,\langle\psi^{j,\theta}|\phi^+_{ab}\rangle\;\;\mathrm{Tr}\big[\sigma_{\mathsf{R}}^{i,j,\theta}\big] \qquad\qquad \text{(by def.)}\\
&\leq\;2^{-t-1}\sum_{i,j\in I^\theta} q_{ij}\sum_{a,b\in\{0,1\}^m}\langle\phi^+_{ab}|\psi^{i,\theta}\rangle\,\langle\psi^{j,\theta}|\phi^+_{ab}\rangle\;\mathrm{Tr}\big[\sigma_{\mathsf{R}}^{i,j,\theta}\big] \qquad\qquad \text{(Eq. (3.29))}\\
&=\;2^{-t-1}\sum_{i,j\in I^\theta} q_{ij}\mathrm{Tr}\big[\,|\psi^{j,\theta}\rangle\langle\psi^{j,\theta}|_{\mathsf{XY}}\,\big]\,\mathrm{Tr}\big[\sigma_{\mathsf{R}}^{i,j,\theta}\big]\quad=\quad 2^{-t-1}\,\mathrm{Tr}\big[\varrho_\theta\big]\quad=\quad 2^{-t-1}.
\end{aligned}$$

In the last line, we used that $\big\{\,|\phi^+_{ab}\rangle : a,b\in\{0,1\}^m\big\}$ is an orthogonal basis for XY. Thus, we get

$$\mathbb{E}_\theta\mathrm{Tr}\Big[\big(\mathbb{1}-\Pi_t^{\mathsf{EPR}}\big)\otimes\mathbb{1}_{\mathsf{R}}\,\varrho_\theta\Big]\;\leq\;2^{-t-1}.$$

Plugging this bound in (3.28) and then into (3.26) gives

$$\Pr\big[v'=v\;\wedge\;\mathsf{ok}=1\big]\leq 2^{-m}(m+1)^{2t}+2^{\frac{-t-1}{2}}\,. \tag{3.32}$$

Choosing $t\approx\sqrt{m}$ makes the RHS negligible. □

Corollary 1. $p(\widetilde{H}_0)=\mathsf{negl}(\lambda)$.

As we argued earlier, this concludes the proof of Lemma (13), and thus of Theorem 5.

Extension to compute-and-compare programs

In this section, we show that an SSL scheme for point functions, which is secure with respect to the appropriate program and challenge ensembles, implies an SSL scheme for compute-and-compare programs with the same level of security.

The idea is simple: to lease the compute-and-compare program $\mathsf{CC}[f,y]$, we lease a a program for the point function P_y, and hand out f in the clear. By leasing P_y we are protecting the portion of the compute-and-compare program which checks equality with y. The intuition is that this is sufficient to make the functionality unclonable since its output is not already determined by f. More generally, one might conjecture that, to obtain an SSL scheme for the function $F=f_1\circ f_2\ldots\circ f_\ell$, it is sufficient to lease any of the functions f_i that is sufficiently non-constant *within its context*.

Let ($\mathsf{SSL\text{-}PF.Gen}$, $\mathsf{SSL\text{-}PF.Lease}$, $\mathsf{SSL\text{-}PF.Eval}$, $\mathsf{SSL\text{-}PF.Verify}$) be any SSL scheme for point functions. The compute-and-compare program scheme is defined as follows:

Construction 4 (SSL scheme for compute-and-compare programs)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. The secure software leasing scheme* $\mathsf{SSL\text{-}CC} = (\mathsf{Gen}, \mathsf{Lease}, \mathsf{Eval}, \mathsf{Verify})$ *for compute-and-compare programs is defined by the following* QPT *algorithms:*

- $\mathsf{SSL\text{-}CC.Gen}(1^\lambda)$*: Takes as input the security parameter λ. Then,*
  - *Let* $\mathsf{sk} \leftarrow \mathsf{SSL\text{-}PF.Gen}(1^\lambda)$. *Output* sk.
- $\mathsf{SSL\text{-}CC.Lease}(1^\lambda, \mathsf{sk}, (f, y))$*: Takes as input a security parameter λ, a secret key sk, and a compute-and-compare program* $\mathsf{CC}[f, y]$*, specified succinctly by f and y. Then,*
  - *Let* $\varrho = \mathsf{SSL\text{-}PF.Lease}(1^\lambda, \mathsf{sk}, y))$. *Output* (f, ϱ).
- $\mathsf{SSL\text{-}CC.Eval}(1^\lambda, (f, \varrho); x)$*: Takes as input a security parameter λ, an alleged program copy (f, ϱ), and a string $x \in \{0, 1\}^n$ (where n is the size of the inputs to f). Then,*
  - *Compute* $y' = f(x)$.
  - *Let* $b \leftarrow \mathsf{SSL\text{-}PF.Eval}(\varrho; y')$. *Output* b.
- $\mathsf{SSL\text{-}CC.Verify}(1^\lambda, \mathsf{sk}, (f, \varrho); \sigma)$*:*
  - *Let* $b' \leftarrow \mathsf{SSL\text{-}PF.Verify}(1^\lambda, \mathsf{sk}, y; \sigma)$. *Output* b'.

Before we state the theorem, we first introduce several classes of distributions over compute-and-compare programs and input challenges. First, we define the distribution $\mathcal{D}_{\mathsf{CC\text{-}UNP}}$ as follows.

- $\mathcal{D}_{\mathsf{CC\text{-}UNP}}$. We refer to this class as the class of *unpredictable compute-and-compare programs*. This consists of ensembles $D = \{D_\lambda\}$ where D_λ is a distribution over compute-and-compare programs such that $\mathsf{CC}[f, y] \leftarrow D_\lambda$ satisfies $\mathbf{H}_{\min}(y|f) \geq \lambda^\epsilon$ for some $\epsilon > 0$, and where the input length of f is λ and the output length is bounded by some polynomial $t(\lambda)$.

We also define the following class of distributions over input challenges:

- $\mathcal{D}_{\mathsf{CC\text{-}Chall\text{-}SSL}}$. An ensemble $D = \{D_{f,y}\}$, where each $D_{f,y}$ is a distribution over the domain of f, belongs to the class $\mathcal{D}_{\mathsf{CC\text{-}Chall\text{-}SSL}}$ if there exists an efficiently sampleable family $\{X_\lambda\}$ of distributions over $\{0, 1\}^\lambda$ with $\mathbf{H}_{\min}(X_\lambda) \geq \lambda^\epsilon$, for some $\epsilon > 0$, and an efficiently sampleable family $\{Z_{f,y}\}$, where $Z_{f,y}$ is a distribution over the set $f^{-1}(y)$, such that $D_{f,y}$ is the following distribution (where λ is the size of inputs to f):
  - with probability $1/2$, sample $z \leftarrow Z_{f,y}$ and output z.

- with probability $1/2$, sample $x \leftarrow X_\lambda$, and output x.

We say the ensemble D is *specified* by the families $\{X_\lambda\}$ and $\{Z_{f,y}\}$.

We also define two classes of distributions over pairs of programs and challenges.

- $\mathcal{D}_{\textsf{CC-pairs-stat-SSL}}$. This consists of pairs of ensembles $\left(D = \{D_\lambda\}, D' = \{D'_{f,y}\}\right)$ where $D \in \mathcal{D}_{\textsf{CC-UNP}}$ and $D' \in \mathcal{D}_{\textsf{CC-Chall-SSL}}$ satisfying the following. Let D' be specified by the families $\{X_\lambda\}$ and $\{Z_{f,y}\}$, and denote by $\textsf{MarkedInput}\left(D_\lambda, \{Z_{f,y}\}\right)$ the distribution over $\{0,1\}^\lambda$ induced by D_λ and $\{Z_{f,y}\}$, i.e.:

  - Sample $(f, y) \leftarrow D_\lambda$, then output $z \leftarrow Z_{f,y}$.

  For any fixed f_* with domain $\{0,1\}^\lambda$ such that (f_*, y_*) is in the support of D_λ for some y_*, denote by $\textsf{MarkedInput}(D_\lambda, \{Z_{f,y}\})|_{f_*}$, the distribution $\textsf{MarkedInput}(D_\lambda, \{Z_{f,y}\})$ conditioned on D_λ sampling f_*. Then, we require that, for any sequence $\{f_*^{(\lambda)}\}$ (where, for all λ, $(f_*^{(\lambda)}, y_*)$ is in the support of D_λ for some y_*) the families $\{X_\lambda\}$ and $\{\textsf{MarkedInput}(D_\lambda, \{Z_{f,y}\})|_{f_*^{(\lambda)}}\}$ are statistically indistinguishable.

- $\mathcal{D}_{\textsf{CC-pairs-comp-SSL}}$. This is defined in the same way as $\mathcal{D}_{\textsf{CC-pairs-stat-SSL}}$, except that we only require $\{X_\lambda\}$ and $\{\textsf{MarkedInput}(D_\lambda, \{Z_{f,y}\})|_{f_*^{(\lambda)}}\}$ to be *computationally* indistinguishable.

Theorem 6. *Let* ($\textsf{SSL-PF.Gen}$, $\textsf{SSL-PF.Lease}$, $\textsf{SSL-PF.Eval}$, $\textsf{SSL-PF.Verify}$) *be an* $\textsf{SSL}$ *scheme for point functions that is secure with respect to all pairs* $(\mathcal{D}, \mathcal{D}') \in \mathcal{D}_{\textsf{PF-pairs-stat-SSL}}$ ($\in \mathcal{D}_{\textsf{PF-pairs-comp-SSL}}$)*. Then, the scheme of Construction 4 is a secure* $\textsf{SSL}$ *scheme for compute-and-compare programs with respect to all pairs* $(D, D') \in \mathcal{D}_{\textsf{CC-pairs-stat-SSL}}$ ($\in \mathcal{D}_{\textsf{CC-pairs-comp-SSL}}$)*. The same conclusion holds relative to any oracle, i.e., when all algorithms have access to the same oracle, with respect to query-bounded (computationally bounded) adversaries.*

We now give a proof of Theorem 6 via a reduction to the point function security experiment.

Proof of Theorem 6. We prove the claim for $\left(\{D_\lambda\}, \{D_{f,y}\}\right) \in \mathcal{D}_{\textsf{CC-pairs-stat-SSL}}$ only, since the case of $\left(\{D_\lambda\}, \{D_{f,y}\}\right) \in \mathcal{D}_{\textsf{CC-pairs-comp-SSL}}$ is virtually identical. Let $t(\lambda)$ be the length of strings in the range of f's sampled from D_λ and let the ensemble $\{D_{f,y}\}$ be specified by $\{X_\lambda\}$ and $\{Z_{f,y}\}$ (using the notation introduced above for ensembles in $\mathcal{D}_{\textsf{CC-Chall-SSL}}$).

Let $\mathcal{A}$ be an adversary for the compute-and-compare $\textsf{SSL}$ scheme of Construction 4 with respect to ensembles $\{D_\lambda\}$ and $\{D_{f,y}\}$ who wins at the $\textsf{SSL}$ security game with probability $p(\lambda) > 0$. It then follows that for each λ there exists $f_*^{(\lambda)}$ such that $(f_*^{(\lambda)}, y)$ is in the support of D_λ for some y, and such that the probability that $\mathcal{A}$ wins is at least $p(\lambda)$, conditioned on $f_*^{(\lambda)}$ being sampled.

We will construct an adversary $\mathcal{A}'$ that wins with probability $p(\lambda) - \mathsf{negl}(\lambda)$ in the point function security game with respect to the distributions $\{D'_{t(\lambda)}\}$ and $\{D'_y\}$, defined as follows:

- $D'_{t(\lambda)}$: sample $x \leftarrow X_\lambda$ and output the point function $P_{f_*^{(\lambda)}(x)}$.
- D'_y: sample $x \leftarrow D_{f_*^{(\lambda)},y}$ and output $f_*^{(\lambda)}(x)$.

The adversary $\mathcal{A}'$ against the point function SSL game acts as follows:

- $\mathcal{A}'$ receives a state ϱ from the lessor, and then forwards $(f_*^{(\lambda)}, \varrho)$ to adversary $\mathcal{A}$.
- $\mathcal{A}$ returns a supposed program copy σ for the point function to $\mathcal{A}'$ who then sends it back to the lessor for verification.
- Conditioning on the verification being successful, the lessor replies with a challenge input $x \leftarrow D'_y$. $\mathcal{A}'$ then samples $x' \leftarrow Z_{f,x}$, and runs $\mathcal{A}$ with input challenge x'.
- Let b be the bit returned by $\mathcal{A}$. The adversary $\mathcal{A}'$ replies with the same b to the lessor.

It is straightforward to check that the game "simulated" by $\mathcal{A}'$ for $\mathcal{A}$ is statistically indistinguishable from a security game with respect to $\{D_\lambda\}$ and $\{D_{f,y}\}$, conditioned on $f_*^{(\lambda)}$. Thus, we deduce, by hypothesis, that $\mathcal{A}$ passes verification and returns the correct bit with probability at least $p(\lambda) - \mathsf{negl}(\lambda)$, and thus $\mathcal{A}'$ wins with probability at least $p(\lambda) - \mathsf{negl}(\lambda)$. Crucially, note that $\left(\{D'_{t(\lambda)}\}, \{D'_y\}\right) \in \mathcal{D}_{\mathsf{PF\text{-}pairs\text{-}stat\text{-}SSL}}$. It follows that if the SSL is secure, then the compute-and-compare scheme must also be secure.

We remark that the proof of the theorem statement relative to any oracle is analogous. $\square$

Chapter 4

REVOKING ENCRYPTED DATA: PUBLICLY-VERIFIABLE DELETION

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. Broadbent and Islam [40] recently built 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 quantum ciphertext has been deleted by an untrusted party, it offers no additional layer of functionality.

In this chapter, we use Gaussian superpositions over lattices to construct the first fully homomorphic encryption scheme with certified deletion—a protocol that enables an untrusted quantum server to compute on encrypted data and to also prove data deletion to a client.

Organization. First, we prove some basic facts about Gaussian superpositions in Section 4.2. Next, in Section 4.3, we generalize the notion of *collapsing* hashes and prove the strong Gaussian-collapsing property of the Ajtai hash function; this marks the main technical result of this chapter. In Section 4.4 we define the syntax and security of public-key encryption schemes with publicly-verifiable deletion. Then, in Section 4.5, we construct a Dual-Regev public-key encryption scheme with publicly-verifiable deletion and prove its security. In Section 4.6, we define the syntax and security of homomorphic encryption schemes with publicly-verifiable deletion, which is analogous as in the public-key setting. In Section 4.7, we give the main construction of this chapter; namely, our Dual-Regev (leveled) fully homomorphic encryption scheme with publicly-verifiable deletion. We prove that it achieves certified deletion security using a similar proof as for our public-key scheme. Finally, in Section 4.8, we describe a four-message protocol for FHE with simultaneous data deletion, which allows an untrusted quantum server to compute on encrypted data and to simultaneously prove data deletion to a client – all in a single interactive protocol.

4.1 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* [63]. While new data privacy regulations have been put into practice in several jurisdictions, formalizing data deletion remains a fundamental challenge for classical 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* [108, 59] can achieve a limited notion of data deletion under *bounded memory assumptions*. Recently, Garg, Goldwasser and Vasudevan [63] 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. [63] 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 [51, 41]. Contrary to classical data, it is fundamentally impossible to create copies of an unknown quantum state thanks to the *quantum no-cloning theorem* [135]. Broadbent and Islam [41] proposed a quantum encryption scheme which enables a user to certify the deletion of a quantum ciphertext. Unlike classical proofs-of-secure-erasure, this cryptographic notion of certified deletion is achievable unconditionally in a fully malicious adversarial setting [41]. 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 [41] 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 homomorphic encryption schemes [113, 66, 37] 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 [41] requires a classical *one-time pad*, whereas the authors in [83] 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 [35], 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 a malicious adversary to break the encryption scheme retrospectively – long after the execution of the protocol. Long-term security is especially relevant for data which is required to remain confidential for many years; for example, such as private medical records or sensitive 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 also enabling the server to prove data deletion to a client, thus effectively achieving a form of *everlasting security* [104, 82].

Technical overview

How can we certify that sensitive information has been deleted by an untrusted party? Quantum information allows us to achieve a cryptographic notion called *certified deletion* [51, 61, 41]. The main idea behind this concept is the *principle of complementarity*. This feature allows us to encode information in mutually incompatible bases—a notion that has no counterpart in a classical world.

Broadbent and Islam [41] construct a private-key quantum encryption scheme with certified deletion using a BB84-type protocol that closely resembles the standard quantum key distribution (QKD) protocol [29, 122]. The crucial idea behind the scheme is that the information which is necessary to decrypt is encoded in the *computational basis*, whereas *certifying deletion* requires a measurement in the incompatible *Hadamard basis*. The scheme in [41] achieves a rigorous notion of certified deletion security: once the ciphertext is successfully deleted, the plaintext m remains hidden even if the private key is later revealed. Using a standard *hybrid encryption scheme*, Hiroka, Morimae, Nishimaki and Yamakawa [83] extended the scheme in [41] to both public-key and attribute-based encryption with certified deletion via the notion of *receiver non-committing* (RNC) encryption [87, 45]. The security proof in[83] 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 ciphertexts with no additional layer of functionality. In this chapter, we answer a question raised by Broadbent and Islam [41] 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 insufficient in order to construct homomorphic encryption with publicly-verifiable deletion, and thus an entirely new approach is necessary.

Our techniques deviate from the hybrid encryption paradigm of previous works [41, 82] and allow us to construct the *first* homomorphic quantum encryption scheme with certified deletion which has the desirable feature of being publicly verifiable. The main technical ingredient of our scheme is a protocol by which a quantum prover can convince a classical verifier that a sample from the *Learning with Errors* [112] distribution in the form of a quantum state was deleted.

Quantum superpositions of LWE samples. The *Learning with Errors* (LWE) problem was introduced by Regev [112] and has given rise to numerous cryptographic applications, including public-key encryption [68], homomorphic encryption [37, 69] and attribute-based encryption [33].

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 $(\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}, \mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q)$ 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 a uniformly random vector and $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ is a vector which is sampled according to the discrete Gaussian distribution $D_{\mathbb{Z}^m,\alpha q}$. The latter 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 whether a malicious party has deleted a sample from the LWE distribution? Our main technical insight 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 [77] in the context of quantum learning theory and by Alagic, Jeffery, Ozols, and Poremba [11], as well as by Chen, Liu, and Zhandry [48], in the context of quantum cryptanalysis of LWE-based cryptosystems. Let us now describe the main idea behind our constructions. Consider the Gaussian superposition,[1]

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

Here, we let $\sigma = 1/\alpha$ and 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

$$|\hat{\psi}_\mathbf{y}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x})\, |\mathbf{x}\rangle . \tag{4.1}$$

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 12). The quantum state $|\hat{\psi}_\mathbf{y}\rangle$ in Eq. (4.1) has the

[1] A tail bound shows that $D_{\mathbb{Z}^m,\sigma}$ is essentially only supported on $\{\mathbf{x} \in \mathbb{Z}^m : \|\mathbf{x}\| \leq \sigma\sqrt{m}\}$. 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.

following *duality property*; namely, by applying the (inverse) q-ary quantum Fourier transform, we obtain the state

$$|\psi_{\mathbf{y}}\rangle = \sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m} \varrho_{\frac{q}{\sigma}}(\mathbf{e})\,\omega_q^{-\langle\mathbf{s},\mathbf{y}\rangle}\,|\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal \pmod{q}\rangle\,, \tag{4.2}$$

where $\omega_q = e^{2\pi i/q}$ is the primitive q-th root of unity. We make this statement more precise in Lemma 20. 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. [120] who gave quantum reduction from SIS to LWE based on Regev's reduction [112], and was later implicitly used by Roberts [115] and Kitagawa et al. [90] to construct quantum money and secure software leasing schemes.

Our quantum encryption schemes with certified deletion exploit the fact that a 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 ciphertext information which is necessary to decrypt is encoded using LWE samples in the *computational basis*.

Dual-Regev public-key encryption with publicly-verifiable deletion. The key ingredient of our homomorphic encryption scheme with certified deletion is the *Dual-Regev* public-key encryption scheme introduced by Gentry, Peikert, and Vaikuntanathan [68]. Using Gaussian states, we can encode Dual-Regev ciphertexts for the purpose of certified deletion while simultaneously preserving their full cryptographic functionality. Our scheme consists of the following efficient algorithms:

- To generate a pair of keys $(\mathsf{sk},\mathsf{pk})$, sample a random $\mathbf{A}\in\mathbb{Z}_q^{n\times(m+1)}$ together with a particular short trapdoor vector $\mathbf{t}\in\mathbb{Z}^{m+1}$ such that $\mathbf{A}\cdot\mathbf{t}=\mathbf{0} \pmod{q}$. Let $\mathsf{pk}=\mathbf{A}$ and $\mathsf{sk}=\mathbf{t}$.

- To encrypt $b\in\{0,1\}$ using $\mathsf{pk}=\mathbf{A}$, generate the following for a random $\mathbf{y}\in\mathbb{Z}_q^n$:

$$\mathsf{vk}\leftarrow(\mathbf{A},\mathbf{y}),\qquad |\mathsf{CT}\rangle\leftarrow\sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^{m+1}}\varrho_{q/\sigma}(\mathbf{e})\,\omega_q^{-\langle\mathbf{s},\mathbf{y}\rangle}\,|\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal+b\cdot(0,\dots,0,\lfloor\frac{q}{2}\rfloor)\rangle\,,$$

  where vk is a public verification key and $|\mathsf{CT}\rangle$ is the quantum ciphertext for $\sigma>0$.

- To decrypt $|\mathsf{CT}\rangle$ using sk, measure in the computational basis to obtain $\mathbf{c}\in\mathbb{Z}_q^{m+1}$, and output 0, if $\mathbf{c}^\intercal\cdot\mathsf{sk}\in\mathbb{Z}_q$ is closer to 0 than to $\lfloor\frac{q}{2}\rfloor$, and output 1, otherwise. Here $\mathsf{sk}=\mathbf{t}$ is chosen such that $\mathbf{c}^\intercal\cdot\mathsf{sk}$ yields an approximation of $b\cdot\lfloor\frac{q}{2}\rfloor$ from which we can recover b.

To delete the ciphertext $|\mathsf{CT}\rangle$, we simply perform measurement in the Fourier basis. In Corollary 2, we show that the Fourier transform of the ciphertext $|\mathsf{CT}\rangle$ results in the *dual* quantum state

$$|\widehat{\mathsf{CT}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^{m+1}:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x})\, \omega_q^{\langle \mathbf{x}, b\cdot(0,\dots,0,\lfloor \frac{q}{2}\rfloor)\rangle}\, |\mathbf{x}\rangle\,. \tag{4.3}$$

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 [103, 68], it is computationally difficult to produce a valid deletion certificate from $(\mathbf{A}, \mathbf{y})$ alone.

To formalize security, we consider the notion of (everlasting) *certified deletion security* (i.e., $\mathsf{EV\text{-}CD}$ security) which was proposed by Bartusek and Khurana [23] as a strengthening of the original notion by Broadbent and Islam [40]. Roughly speaking, $\mathsf{EV\text{-}CD}$ security guarantees that, once deletion of the ciphertext is successful, the plaintext remains hidden even if the adversary is subsequently allowed to run in unbounded time (see Definition 34). We prove the security of our schemes by exploiting a strong *collapsing*-type property of the Ajtai hash function which we show under the quantum hardness of LWE and SIS. This is our main technical result in this chapter.

Gaussian-collapsing hash functions. Unruh [125] 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. (4.1), 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 30), 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 [97] implicitly showed that the *Ajtai 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* and (decisional) LWE. As a first

$$\begin{tikzcd}
\displaystyle\sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x})\ |\mathbf{x}\rangle \arrow[r, phantom, "\approx_c", "(\text{Thm. 8})"'] \arrow[d, leftrightarrow, "\mathsf{FT}_q"', "(\text{Lem. 20})"] & |\mathbf{x}_0\rangle,\quad \mathbf{x}_0 \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}),\frac{\sigma}{\sqrt{2}}} \arrow[d, leftrightarrow, "\mathsf{FT}_q"] \\
\displaystyle\sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m} \varrho_{\frac{q}{\sigma}}(\mathbf{e})\,\omega_q^{-\langle\mathbf{s},\mathbf{y}\rangle}|\mathbf{s}^\intercal\mathbf{A}+\mathbf{e}^\intercal\rangle \arrow[r, phantom, "(\text{Thm. 9})", "\approx_c"'] & \displaystyle\sum_{\mathbf{u}\in\mathbb{Z}_q^m} \omega_q^{-\langle\mathbf{u},\mathbf{x}_0\rangle}\ |\mathbf{u}\rangle
\end{tikzcd}$$

Figure 4.1: Technical overview of Gaussian superposition states and their properties used throughout this work. The computational indistinguishability property holds under the (subexponential) hardness of the LWE assumption (Definition 14). Here, $\Lambda_q^{\mathbf{y}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\cdot\mathbf{x} = \mathbf{y} \pmod q\}$ denotes a coset of the lattice $\Lambda_q^{\perp}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\cdot\mathbf{x} = \mathbf{0} \pmod q\}$ defined in Section 2.6.

preliminary result, we give a simple and direct proof (Theorem 8) of the Gaussian-collapsing property assuming the hardness of decisional LWE, which might be of independent interest.

The fact Ajtai's hash function is Gaussian-collapsing has several implications for the security of our schemes. Because our Dual-Regev ciphertext corresponds to the Fourier transform of the state in Eq. (4.3), the Gaussian-collapsing property immediately implies the semantic (i.e., IND-CPA) security under decisional LWE (see Theorem 9). We refer to Figure 4.1 for an overview of our Gaussian states and their properties.

To prove the stronger notion of EV-CD security of our Dual-Regev scheme with publicly-verifiable deletion, we have to show that, once deletion has taken place, the plaintext remains information-theoretically hidden from the view of the adversary. We observe that it is sufficient to show that Ajtai's hash function satisfies a particular *strong Gaussian-collapsing property*; namely, once an 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 it received a Gaussian superposition of pre-images or a single (measured) pre-image, even if $\mathcal{A}$ is now allowed to run in unbounded time. Here, it is crucial that $\mathcal{A}$ is unbounded 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 distinguishing between a superposition of LWE samples and a uniform superposition. We prove the strong Gaussian-collapsing property in Theorem 10, assuming the hardness of LWE and SIS.

We then go on to prove the following result in Theorem 12:

Theorem (Informal)**.** *Our Dual-Regev encryption scheme with publicly-verifiable deletion (see Construction 5) is* $\mathsf{EV\text{-}CD}$*-secure, assuming the quantum subexponential hardness of* LWE *and* SIS.

To gain some intuition for why the strong Gaussian-collapsing property holds, 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[2] 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 succeeded without disturbing the initial quantum state, we could potentially provide a valid certificate π and also distinguish the auxiliary system once we are allowed to be computationally unbounded. However, by shifting the state by another LWE sample, we have necessarily entangled the two systems in a way that prevents us from finding a valid certificate via a Fourier basis measurement. We give a formal proof of the strong Gaussian-collapsing property in Theorem 10.

Next, we extend our Dual-Regev scheme towards a (leveled) FHE scheme with certified deletion.

Dual-Regev fully homomorphic encryption with publicly-verifiable deletion. Our (leveled) FHE scheme with certified deletion is based on the (classical) Dual-Regev leveled FHE scheme used by Mahadev [99]—a variant of the scheme due to Gentry, Sahai and Waters [69]. 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 $N = (m + 1)\lceil \log q \rceil$ and let $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ denote the so-called power-of-2 *gadget matrix* (defined in Section 4.7). The scheme consists of the following efficient algorithms:

- To generate a pair of keys $(\mathsf{sk}, \mathsf{pk})$, sample $\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}$ together with a particular short trapdoor vector $\mathbf{t} \in \mathbb{Z}^{m+1}$ 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 $\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}$, generate the following pair consisting of a classical verification key and quantum ciphertext for a random $\mathbf{Y} \in \mathbb{Z}_q^{n\times N}$ with columns $\mathbf{y}_1, \ldots, \mathbf{y}_N \in \mathbb{Z}_q^n$:

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

  where $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ denotes the *gadget matrix* and where $\sigma = 1/\alpha$.

[2] 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.

- 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.

Deletion and verification take place exactly as in our Dual-Regev scheme with certified deletion.

In Theorem 15, we prove that our scheme satisfies the same notion of certified deletion security which we previously considered in the context of (regular) public-key encryption.

Theorem. *Our Dual-Regev (leveled)* FHE *scheme with publicly-verifiable deletion (Construction 7) is* EV-CD*-secure, assuming the quantum subexponential hardness of* LWE *and* SIS*.*

Our FHE scheme supports the evaluation of polynomial-sized Boolean circuits consisting entirely of NAND gates, which are universal for classical computation. Inspired by the classical homomorphic NAND operation of the Dual-Regev scheme [69, 99], we define an analogous quantum operation U_{NAND} in Definition 38 which allows us to apply a NAND gate directly onto Gaussian states. When applying homomorphic operations, the new ciphertext 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. Notice, however, that the resulting ciphertext 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 server perform homomorphic computations and, if requested, to also 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 prevents us from obtaining a valid deletion certificate via a Fourier basis measurement. At first sight, it seems as if applying homomorphic operations and proving data deletion are two mutually exclusive properties. Indeed, our basic Dual-Regev encryption scheme in Construction 7 only supports homomorphic operations and publicly-verifiable deletion as separate properties.

It is natural to ask whether it is possible to achieve both tasks simultaneously, say in a protocol between a client and an untrusted server. Remarkably, such a protocol would allow an untrusted server to compute on private data and, if requested, to simultaneously prove data deletion to a client. We show that such a protocol is indeed possible, albeit with a few important caveats which we explain in Section 4.8. In Protocol 1, we describe a four-message protocol for FHE with simultaneous data deletion which is based on our Dual-Regev FHE scheme in Construction 7. To resolve the aforementioned technical issue, we introduce additional interaction between the server and the client (which is not required for a conventional homomorphic encryption scheme). After performing a 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} corresponding to an encryption of $x = (x_1, \dots, x_\ell) \in \{0,1\}^\ell$, the server simply sends the quantum system C_{out} containing an encryption of $C(x)$ to the client. Then, using the secret key sk (i.e., a trapdoor for the public matrix pk), it is possible for the client to *extract* the outcome $C(x)$ from the system C_{out} with overwhelming probability without significantly damaging the state. We show that 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 client has learned the outcome of the homomorphic application of the circuit C while the server is still in possession of a large number of auxiliary systems (denoted by C_{aux}) which mark intermediate applications of the gate U_{NAND}. We remark that this is where the standard FHE protocol ends. In order to enable *certified deletion*, the client must now return the system C_{out} to the server. Having access to all three systems $C_{\mathsf{in}} C_{\mathsf{aux}} C_{\mathsf{out}}$, the server 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 server's possession is now approximately a simple product of Gaussian states, the server can perform a Fourier basis measurement of systems C_{in}, as required. Once the protocol is complete, it is therefore possible for the client to know $C(x)$ and to be convinced that data deletion has taken place. In Section 4.8, we show that our four-message protocol indeed achieves certified deletion, provided that the server is *honest* during the homomorphic evaluation phase of the protocol.

Related work

The first work to formalize a notion resembling *certified deletion* is due to Unruh [127] 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. Fu and Miller [61] gave the first quantum protocol that proves deletion of a single bit using classical interaction alone. Subsequently, Coiteux-Roy and Wolf [51] proposed a QKD-like conjugate coding protocol that enables certified deletion of a classical plaintext, albeit without a complete security proof. Independently of [51], Broadbent and Islam [41] construct a private-key quantum encryption scheme with a rigorous definition of certified deletion using a BB84-type protocol that closely resembles the standard quantum key distribution protocol [29, 122]. 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 *certifying 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 private-key quantum encryption scheme of Broadbent and Islam [41] 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 [83] extended the scheme in [41] to both public-key and attribute-based encryption with certified deletion via the notion of *receiver non-committing* (RNC) encryption [87, 45]; for example, to obtain a public-key encryption scheme with certified deletion, one simply outputs a quantum ciphertext of the [41] scheme together with a classical (non-committing) public-key encryption of its private key. Given access to the RNC secret key, it is therefore possible to decrypt the quantum ciphertext. Crucially, the hybrid encryption scheme also inherits the certified deletion property of the [41] scheme; namely, once deletion has taken place, the plaintext remains hidden even if the RNC secret key is later revealed. The security proof in [83] 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. To obtain a homomorphic encryption scheme with certified deletion, one would have to instantiate the hybrid encryption scheme with a classical (non-committing) homomorphic encryption scheme which is not known to exist. While generic transformations for non-committing encryption have been studied [91], they tend to be incompatible with basic homomorphic computations. Moreover, it is unclear whether the candidate hybrid approach for homomorphic encryption is even secure: for all we know, a malicious adversary could use homomorphic evaluation to decouple the quantum part from the classical part of the ciphertext in order to obtain a classical encryption of the plaintext, thereby violating certified deletion security.

Hiroka, Morimae, Nishimaki and Yamakawa [82] studied *certified everlasting zero-knowledge proofs* for QMA via the notion of *everlasting security* which was first formalized by Müller-Quade and Unruh [104]. A recent paper by Coladangelo, Liu, Liu, and Zhandry [53] introduces *subspace coset states* in the context of unclonable crytography in a way that loosely resembles our use of primal and dual Gaussian states. In concurrent work, Bartusek and Khurana [23] consider generic transformations for encryption schemes with certified deletion. Similar to Broadbent and Islam [41], they use a hybrid approach via BB84 states to construct public-key, attribute-based and homomorphic encryption schemes with *certified everlasting security*: once deletion is successful, the security notion guarantees that the plaintext remains hidden even if the adversary is henceforth computationally unbounded. However, in contrast with our results, their notion only considers deletion certificates which are privately verifiable. Finally, we remark that subsequent follow-up work by Bartusek et al. [25] has since constructed encryption schemes with public verification (assuming the existence of post-quantum indistinguishability obfuscation) as well as *maliciously* secure bind delegation protocol with certified deletion via succinct non-interactive arguments (SNARGs) for polynomial-time computation, which can be constructed from LWE.

4.2 Primal and Dual Gaussian States

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. [120] who gave a quantum reduction from the SIS problem to the LWE problem. Given $q \in \mathbb{N}$, $m \in \mathbb{N}$ and $\sqrt{8m} < \sigma < q/\sqrt{8m}$, a Gaussian superposition over $\mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$ is a pure state of the form

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

Note that the state $|\psi\rangle$ is not normalized for convenience and ease of notation. The tail bound in Lemma 11 implies that (the normalized variant of) $|\psi\rangle$ is within negligible trace distance of a *truncated* discrete Gaussian superposition $|\tilde{\psi}\rangle$ with support $\{\mathbf{x} \in \mathbb{Z}_q^m : \|\mathbf{x}\| \leq \sigma\sqrt{\frac{m}{2}}\}$, where

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

In this section, we consider Gaussian superpositions with parameter $\sigma = \Omega(\sqrt{m})$ which can be efficiently implemented using standard quantum state preparation techniques; for example using *quantum rejection sampling* and the *Grover-Rudolph algorithm* [78, 112, 36, 38].

Our Dual-Regev-type encryption schemes with certified deletion in Section 4.5 and Section 4.7 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 2.6.

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 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.

We define primal and dual Gaussian states as follows.

Definition 28 (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*

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

- *(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.

Duality lemma

In this section, we prove a lemma which states that, up to negligible trace distance, the primal and dual Gaussian states in Definition 28 are related via the q-ary quantum Fourier transform.

First, we show the following technical result.

Lemma 19. *Let $m \in \mathbb{N}$, $q \geq 2$ a modulus and let $\sigma > 0$. Consider the quantum states,*

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

where $\varrho_{\sigma,q}$ is the periodic Gaussian from Definition 7. Then, the normalized variants of the Gaussian superpositions $|\psi\rangle$ and $|\phi\rangle$ above satisfy

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

Proof. We consider the following two distributions over $\mathbb{Z}^m \cap (-\frac{q}{2},\frac{q}{2}]^m$ given by

$$D_{\sigma}(\mathbf{x}) = \frac{\varrho_{\sigma}(\mathbf{x})}{\sum_{\mathbf{y}\in\mathbb{Z}_q^m}\varrho_{\sigma}(\mathbf{y})} \qquad \text{and} \qquad D_{\sigma,q}(\mathbf{x}) = \frac{\varrho_{\sigma,q}(\mathbf{x})}{\sum_{\mathbf{y}\in\mathbb{Z}_q^m}\varrho_{\sigma,q}(\mathbf{y})}. \tag{4.4}$$

We first bound the Hellinger distance,

$$H^2(D_\sigma, D_{\sigma,q}) = 1 - \sqrt{Z_\sigma^{-1} \cdot Z_{\sigma,q}^{-1}} \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \sqrt{\varrho_\sigma(\mathbf{x}) \cdot \varrho_{\sigma,q}(\mathbf{x})}, \tag{4.5}$$

where we define the two normalization factors

$$Z_\sigma = \sum_{\mathbf{y} \in \mathbb{Z}_q^m} \varrho_\sigma(\mathbf{y}) \qquad \text{and} \qquad Z_{\sigma,q} = \sum_{\mathbf{y} \in \mathbb{Z}_q^m} \varrho_{\sigma,q}(\mathbf{y}). \tag{4.6}$$

From Lemma 14, it follows for any $\mathbf{x} \in \mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$ that

$$\varrho_{\sigma,q}^2(\mathbf{x}) \cdot \left(1 + 2^{-(\frac{1}{2}(q/\sigma)^2 - m)}\right)^{-1} \leq \varrho_\sigma(\mathbf{x}) \cdot \varrho_{\sigma,q}(\mathbf{x}). \tag{4.7}$$

Plugging in Eq. (4.7), we can bound the Hellinger distance as follows:

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

Therefore, it holds that

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

□

We are now ready to prove the so-called duality lemma.

Lemma 20 (Duality lemma). *Let $m \in \mathbb{N}$ and $q \geq 2$ be a prime modulus and let $q/\sqrt{8m} > \sigma > \sqrt{8m}$. 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 an arbitrary vector. Then, up to negligible trace distance, the (normalized variants of the) primal and dual Gaussian states are related via the quantum Fourier transform:*

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

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

$$|\hat{\psi}_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} = \mathbf{y} \pmod q}} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \, . \tag{4.8}$$

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

$$|\hat{\phi}_{\mathbf{y}}\rangle = \sum_{\mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A})} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \, . \tag{4.9}$$

Since $\sigma < q/\sqrt{8m}$, it follows from the tail bound in Lemma 12 that the state in (4.8) is within negligible trace distance of the state in Eq. (4.9). Applying the (inverse) Fourier transform, we get

$$|\phi_{\mathbf{y}}\rangle \overset{\text{def}}{=} \mathsf{FT}_q^\dagger \, |\hat{\phi}_{\mathbf{y}}\rangle = \sum_{\mathbf{z} \in \mathbb{Z}_q^m} \Big(\sum_{\mathbf{x} \in \Lambda_q^{\mathbf{y}}(\mathbf{A})} \varrho_\sigma(\mathbf{x}) \cdot \omega_q^{-\langle \mathbf{x}, \mathbf{z} \rangle} \Big) \, |\mathbf{z}\rangle \, . \tag{4.10}$$

From the Poisson summation formula (Lemma 13) and a subsequent change of variables, we get

$$\begin{aligned} |\phi_{\mathbf{y}}\rangle &= \sum_{\mathbf{z} \in \mathbb{Z}_q^m} \Big(\sum_{\mathbf{s} \in \mathbb{Z}_q^n} \varrho_{q/\sigma,q}(\mathbf{z} + \mathbf{A}^\intercal \mathbf{s}) \cdot \omega_q^{\langle \mathbf{s}, \mathbf{y} \rangle} \Big) \, |\mathbf{z}\rangle \\ &= \sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{\mathbf{e} \in Z_q^m} \varrho_{q/\sigma,q}(\mathbf{e}) \cdot \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle} \, |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q\rangle \, . \end{aligned} \tag{4.11}$$

Because $\sigma > \sqrt{8m}$ it follows from Lemma 19 that there exists

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

such that

$$|\phi_{\mathbf{y}}\rangle \approx_\kappa \sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{\mathbf{e} \in Z_q^m} \varrho_{q/\sigma}(\mathbf{e}) \cdot \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle} \, |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod q\rangle \, . \tag{4.12}$$

Putting everything together, it follows from the triangle inequality that

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

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

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

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

Corollary 2. *Let $m \in \mathbb{N}$ and $q \geq 2$ be a prime modulus and let $q/\sqrt{8m} > \sigma > \sqrt{8m}$. 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 an arbitrary vector. Then, there exists a negligible function $\varepsilon : \mathbb{N} \to \mathbb{R}^+$ 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 20 implies that $\mathsf{FT}_q |\psi_{\mathbf{y}}\rangle$ is within negligible trace distance of $|\hat{\psi}_{\mathbf{y}}\rangle$. This proves the claim. □

Efficient state preparation

In this section, we give two algorithms that prepare the *primal* and *dual* Gaussian states from Definition 28. 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 [78, 112, 36, 38]) for a reference.

Our first algorithm (see Algorithm 1 in Figure 4.2) prepares the dual Gaussian state from Definition 28 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 4.3) prepares the primal Gaussian state with respect to an input matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and parameter $\sigma = \Omega(\sqrt{m})$. Here, in order for Lemma 20 to apply, it is crucial that the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$. Fortunately, it follows from Lemma 8 that a uniformly random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ satisfies this property with overwhelming probability.

Invariance under Pauli-Z dephasing

In this section, we prove a surprising property about the dual Gaussian state from Definition 28. We prove Theorem 7, which says that the Pauli-$\mathbf{Z}$ dephasing channel with respect to the LWE distribution leaves the dual Gaussian state approximately invariant.

Theorem 7. *Let $n, m \in \mathbb{N}$ and $q \geq 2$ be a prime modulus, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let σ be a parameter with $q/\sqrt{8m} > \sigma > \sqrt{8m}$. Let $\mathbf{y} \in \mathbb{Z}_q^n$ be arbitrary 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,*

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

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 $1/\alpha = \sigma \cdot 2^{o(n)}$, i.e.,

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

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 $\mathbf{y} \in \mathbb{Z}_q^n$.
Prepare a Gaussian superposition in system $X$ with parameter $\sigma > 0$:
    $|\hat{\psi}\rangle_{XY} = \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle_X \otimes |\mathbf{0}\rangle_Y .$
Apply the unitary $U_{\mathbf{A}} : |\mathbf{x}\rangle \, |\mathbf{0}\rangle \to |\mathbf{x}\rangle \, |\mathbf{A} \cdot \mathbf{x} \ (\mathrm{mod}\ q)\rangle$ on systems $X$ and $Y$:
    $|\hat{\psi}\rangle_{XY} = \sum_{\mathbf{x} \in \mathbb{Z}_q^m} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle_X \otimes |\mathbf{A} \cdot \mathbf{x} \ (\mathrm{mod}\ q)\rangle_Y .$
Measure system $Y$ in the computational basis, resulting in the state
    $|\hat{\psi}_{\mathbf{y}}\rangle_{XY} = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m: \\ \mathbf{A}\mathbf{x} = \mathbf{y}}} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle_X \otimes |\mathbf{y}\rangle_Y .$
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 4.2: 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 28.

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

$$|\hat{\psi}_{\mathbf{y}}\rangle \ = \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} = \mathbf{y} \ (\mathrm{mod}\ q)}} \varrho_\sigma(\mathbf{x}) \, |\mathbf{x}\rangle \, . \tag{4.13}$$

```
Algorithm 2: 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 $\sigma = \Omega(\sqrt{m})$.
Output: Gaussian state $|\psi_{\mathbf{y}}\rangle$ and $\mathbf{y} \in \mathbb{Z}_q^n$.
Run GenDual$(\mathbf{A},\sigma)$, resulting in the state
    $|\hat{\psi}_{\mathbf{y}}\rangle_{XY} = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{Ax}=\mathbf{y}}} \varrho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle_X \otimes |\mathbf{y}\rangle_Y\,.$
Apply the quantum Fourier transform $\mathsf{FT}_q$ to system $X$.
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 4.3: Quantum algorithm which takes as input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and a real parameter $\sigma = \Omega(\sqrt{m})$, and outputs the primal Gaussian state in Definition 28.

Consider a sample $\mathbf{b} = \hat{\mathbf{s}}^\intercal \mathbf{A} + \hat{\mathbf{e}}^\intercal \pmod{q}) \sim \mathsf{LWE}_{n,q,\alpha q}^m$ with $\hat{\mathbf{s}} \xleftarrow{\$} \mathbb{Z}_q^n$ and $\hat{\mathbf{e}} \sim D_{\mathbb{Z}_q^m,\alpha q}$. Because $q/\sqrt{8m} > \sigma > \sqrt{8m}$ and $1/\alpha = \sigma \cdot 2^{o(n)}$, there exist negligible $\eta(\lambda)$ and $\kappa(\lambda)$ such that

$$\begin{aligned}
\mathbf{Z}_q^{\hat{\mathbf{s}}^\intercal\mathbf{A}+\hat{\mathbf{e}}^\intercal}\,|\hat{\psi}_{\mathbf{y}}\rangle &= \mathsf{FT}_q\,\mathbf{X}_q^{\hat{\mathbf{s}}^\intercal\mathbf{A}+\hat{\mathbf{e}}^\intercal}\,\mathsf{FT}_q^\dagger\,|\hat{\psi}_{\mathbf{y}}\rangle && \text{(Lemma 6)}\\
&\approx_\eta \mathsf{FT}_q\,\mathbf{X}_q^{\hat{\mathbf{s}}^\intercal\mathbf{A}+\hat{\mathbf{e}}^\intercal}\,|\psi_{\mathbf{y}}\rangle && \text{(Lemma 20)}\\
&\approx_\kappa \omega_q^{\langle\hat{\mathbf{s}},\mathbf{y}\rangle}\mathsf{FT}_q\,|\psi_{\mathbf{y}}\rangle && \text{(Lemma 15)}\\
&\approx_\eta \omega_q^{\langle\hat{\mathbf{s}},\mathbf{y}\rangle}\,|\hat{\psi}_{\mathbf{y}}\rangle\,. && \text{(Lemma 20)}
\end{aligned}$$

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

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

In other words, $|\hat{\psi}_{\mathbf{y}}\rangle$ in Eq. (4.13) is an approximate eigenvector of the generalized Pauli operator $\mathbf{Z}_q^{\hat{\mathbf{s}}^\intercal\mathbf{A}+\hat{\mathbf{e}}^\intercal}$ with respect to the same matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$. Note that we can simply discard $\omega_q^{\langle\hat{\mathbf{s}},\mathbf{y}\rangle} \in \mathbb{C}$ because it serves as a global phase. Hence, there exists a negligible function $\varepsilon(\lambda)$ such that

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

□

4.3 Gaussian-Collapsing Hash Functions

Unruh [125] 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 29 (Collapsing hash function, [125])**.** *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 $\mathcal{A}$,*

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

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 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 28, 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 30 (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 $\mathcal{A}$,*

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

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.*

Ajtai's hash function

In this section, we give a simple and direct proof that the Ajtai hash function is Gaussian-collapsing assuming (decisional) LWE.

Theorem 8. *Let $n \in \mathbb{N}$ and $q \geq 2$ be a prime with $m \geq 2n \log q$, each parameterized by $\lambda \in \mathbb{N}$. Let σ be a such that $q/\sqrt{8m} > \sigma > \sqrt{8m}$. Then, the Ajtai hash function family $\mathcal{H} = \{H_\lambda\}_{\lambda\in\mathbb{N}}$ with*

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

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 $1/\alpha = \sigma \cdot 2^{o(n)}$.

Proof. Let $\mathcal{A}$ denote the QPT adversary in the experiment $\mathsf{GaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(b)$ for some $b \in \{0, 1\}$. 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 $\mathbf{b} = \hat{\mathbf{s}}^\intercal\mathbf{A} + \hat{\mathbf{e}}^\intercal)$ is either a sample from the LWE distribution with $\hat{\mathbf{s}} \xleftarrow{\$} \mathbb{Z}_q^n$ and $\hat{\mathbf{e}} \sim D_{\mathbb{Z}^m,\alpha q}$, 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

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

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 $\mathsf{LWE}_{n,q,\alpha q}^m$ experiment with advantage at least $1/p(\lambda) - \mathsf{negl}(\lambda)$. We distinguish between the following two cases.

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

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

From Theorem 7 it follows that there exists a negligible function $\varepsilon(\lambda)$ such that

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

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

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

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

Theorem 9. *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 σ be a such that $q/\sqrt{8m} > \sigma > \sqrt{8m}$ and let $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ 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 $1/\alpha = \sigma \cdot 2^{o(n)}$:*

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

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

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

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

Moreover, the distribution of $\mathbf{y} \in \mathbb{Z}_q^n$ is negligibly close 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 the lattice coset of $\Lambda_q^\perp(\mathbf{A})$.

Proof. Let $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ be a random matrix. From Lemma 8 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 9, 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 8 we show that the Ajtai 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 $1/\alpha = \sigma \cdot 2^{o(n)}$. 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,

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

which is the result of a computational basis measurement of $|\hat{\psi}_{\mathbf{y}}\rangle$.[3] Since $q/\sqrt{8m} > \sigma > \sqrt{8m}$, Lemma 12 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 20 we showed that the primal Gaussian state

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

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

□

Strong Gaussian-collapsing property of the Ajtai hash function

In this section, we show our main technical result of this chapter. Specifically, we show that Ajtai's hash function satisfies a particular *strong Gaussian-collapsing property*; namely, once an 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 it received a Gaussian superposition of pre-images or a single (measured) pre-image, even if $\mathcal{A}$ is now allowed to run in unbounded time.

Theorem 10 (Strong (everlasting) Gaussian collapsing property of the Ajtai hash)**.** *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and $m \geq 2n\log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let σ be a such that $q/\sqrt{8m} > \sigma > \sqrt{8m}$ and let $\alpha \in (0,1)$ be a noise ratio such that $1/\alpha = 2^{o(n)}\cdot\sigma$. Let $\mathcal{H} = \{H_\lambda\}_{\lambda\in\mathbb{N}}$ be the Ajtai hash function family with*

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

Then, assuming the hardness of the $\mathsf{LWE}^m_{n,q,\alpha q}$ *and* $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$*, it holds for any pair of adversaries $\mathcal{A} = (\mathcal{A}_0,\mathcal{A}_1)$ consisting of a* QPT *algorithm $\mathcal{A}_0$ and an unbounded algorithm $\mathcal{A}_1$:*

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

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

1. *The challenger first chooses a random hash function $h_{\mathbf{A}} \xleftarrow{\$} H_\lambda$ by sampling $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and then generates a Gaussian superposition state $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}\in\mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A},\sigma)$, with*

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

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

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

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

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

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

$\mathsf{Hyb}_0(b)$: This is the original experiment $\mathsf{StrongGaussCollapseExp}_{\mathcal{H},\mathcal{A},\lambda}(b)$:

1. The challenger samples a random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and then generates a Gaussian superposition state $(|\psi_\mathbf{y}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$, with
$$|\psi_\mathbf{y}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{Ax}=\mathbf{y}}} \varrho_\sigma(\mathbf{x}) \; |\mathbf{x}\rangle \,.$$

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

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

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

$\mathsf{Hyb}_1(b)$: This is the following experiment.

1. The challenger samples a random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and then generates a Gaussian superposition state $(|\psi_\mathbf{y}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$, with
$$|\psi_\mathbf{y}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m \\ \mathbf{Ax}=\mathbf{y}}} \varrho_\sigma(\mathbf{x}) \; |\mathbf{x}\rangle \,.$$

2. The challenger samples a random string, $\mathbf{z} \xleftarrow{\$} \mathbb{Z}_q^m$, prepares a $|+\rangle$ state in system C and applies a controlled-$\mathbf{Z}_q^{\mathbf{z}}$ operation, resulting in the state

$$|\psi_{\mathbf{y}}^{\mathbf{z}}\rangle_{CA} = \frac{1}{\sqrt{2}} \sum_{c \in \{0,1\}} |c\rangle_C \otimes \mathbf{Z}_q^{c \cdot \mathbf{z}} |\psi_{\mathbf{y}}\rangle_X$$

   and sends system X together with $\mathbf{A}, \mathbf{y}$ to the adversary $\mathcal{A}_0$.

3. $\mathcal{A}_0$ replies with a certificate π, and initializes $\mathcal{A}_1$ with its residual (internal) state.

4. The challenger checks if $\boldsymbol{\pi}$ satisfies $\mathbf{A} \cdot \boldsymbol{\pi} = \mathbf{y} \pmod q$ and $\|\boldsymbol{\pi}\| \leq \sigma\sqrt{m/2}$. Then, the challenger measures system C to obtain $c' \in \{0,1\}$ and checks that $c' = b$. If both checks are true, $\mathcal{A}_1$ is run until it outputs a bit b'. Otherwise, $b' \leftarrow \{0,1\}$ is sampled uniformly at random. The output of the experiment is b'.

$\mathsf{Hyb}_2(b)$: This is the following experiment.

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

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

2. The challenger samples a random string, $\mathbf{z} \xleftarrow{\$} \mathbb{Z}_q^m$, prepares a $|+\rangle$ state in system C and applies a controlled-$\mathbf{Z}_q^{\mathbf{z}}$ operation, resulting in the state

$$|\psi_{\mathbf{y}}^{\mathbf{z}}\rangle_{CA} = \frac{1}{\sqrt{2}} \sum_{c \in \{0,1\}} |c\rangle_C \otimes \mathbf{Z}_q^{c \cdot \mathbf{z}} |\psi_{\mathbf{y}}\rangle_X$$

   and sends system X together with $\mathbf{A}, \mathbf{y}$ to the adversary $\mathcal{A}_0$.

3. $\mathcal{A}_0$ replies with a certificate π, and initializes $\mathcal{A}_1$ with its residual (internal) state.

4. The challenger checks if $\boldsymbol{\pi}$ satisfies $\mathbf{A} \cdot \boldsymbol{\pi} = \mathbf{y} \pmod q$ and $\|\boldsymbol{\pi}\| \leq \sigma\sqrt{m/2}$. Then, the challenger applies the projective measurement

$$\left\{ |\psi_\pi^{\mathbf{z}}\rangle\langle\psi_\pi^{\mathbf{z}}|\,, \mathbb{1} - |\psi_\pi^{\mathbf{z}}\rangle\langle\psi_\pi^{\mathbf{z}}| \right\} \qquad \text{with} \qquad |\psi_\pi^{\mathbf{z}}\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \omega_q^{\langle \pi, \mathbf{z}\rangle}\,|1\rangle)$$

   and checks that the first outcome is observed. Finally, the challenger measures system C to obtain $c' \in \{0,1\}$ and checks that $c' = b$. If all three checks are true, $\mathcal{A}_1$ is run until it outputs a bit b'. Otherwise, $b' \leftarrow \{0,1\}$ is sampled uniformly at random. The output of the experiment is b'.

Finally, we also use the following hybrid which is convenient for the sake of the proof.

$\mathsf{Hyb}_2'(b)$: This is the following experiment.

1. The challenger samples a random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and Gaussian vector $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and lets $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}_0 \pmod{q}$.
2. The challenger samples a random string, $\mathbf{z} \xleftarrow{\$} \mathbb{Z}_q^m$, prepares a $|+\rangle$ state in system C and applies a controlled-$\mathbf{Z}_q^{\mathbf{z}}$ operation, resulting in the state
$$|\psi_{\mathbf{y}}^{\mathbf{z}}\rangle_{CA} = \frac{1}{\sqrt{2}} \sum_{c\in\{0,1\}} |c\rangle_C \otimes \mathbf{Z}_q^{c\cdot\mathbf{z}} |\mathbf{x}_0\rangle_X$$
and sends system X together with $\mathbf{A}, \mathbf{y}$ to the adversary $\mathcal{A}_0$.
3. $\mathcal{A}_0$ replies with a certificate π, and initializes $\mathcal{A}_1$ with its residual (internal) state.
4. The challenger checks if $\boldsymbol{\pi}$ satisfies $\mathbf{A} \cdot \boldsymbol{\pi} = \mathbf{y} \pmod{q}$ and $\|\boldsymbol{\pi}\| \leq \sigma\sqrt{m/2}$. Then, the challenger applies the projective measurement
$$\left\{ |\psi_\pi^{\mathbf{z}}\rangle\langle\psi_\pi^{\mathbf{z}}| \, , \mathbb{1} - |\psi_\pi^{\mathbf{z}}\rangle\langle\psi_\pi^{\mathbf{z}}| \right\} \qquad \text{with} \qquad |\psi_\pi^{\mathbf{z}}\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \omega_q^{\langle\pi,\mathbf{z}\rangle} |1\rangle)$$
and checks that the first outcome is observed. Finally, the challenger measures system C to obtain $c' \in \{0,1\}$ and checks that $c' = b$. If all three checks are true, $\mathcal{A}_1$ is run until it outputs a bit b'. Otherwise, $b' \leftarrow \{0,1\}$ is sampled uniformly at random. The output of the experiment is b'.

Before we analyze the probability of distinguishing between the consecutive hybrids, we first show that the following statements hold for the final experiment Hyb_2':

Claim 1. *The following statement holds assuming the quantum hardness of the* $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ *problem: With overwhelming probability, the certificate* π *returned by the adversary in Step 3 in* Hyb_3 *is identical to the pre-image* $\mathbf{x}_0$ *produced by the challenger. In other words,*

$$\Pr\left[\begin{array}{c|c} \mathbf{A}\cdot\boldsymbol{\pi}=\mathbf{y} \pmod{q} \textit{ s.t. } \|\boldsymbol{\pi}\|\leq\sigma\sqrt{m/2} & \mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}} \\ \wedge & \mathbf{y}=\mathbf{A}\cdot\mathbf{x}_0 \pmod{q} \\ \pi \neq \mathbf{x}_0 & \pi\leftarrow\mathcal{A}_0(\mathbf{A},\mathbf{y},|\mathbf{x}_0\rangle\langle\mathbf{x}_0|_X) \end{array} \right] \leq \mathsf{negl}(\lambda),$$

where $|\mathbf{x}_0\rangle\langle\mathbf{x}_0|$ *in system X is the reduced state with respect to the bipartite state*

$$|\psi_{\mathbf{y}}^{\mathbf{z}}\rangle_{CX} = \frac{1}{\sqrt{2}} \sum_{c\in\{0,1\}} |c\rangle_C \otimes \omega_q^{c\cdot\langle\mathbf{x}_0,\mathbf{z}\rangle} |\mathbf{x}_0\rangle_X .$$

Proof. Suppose for the sake of contradiction that the probability is at least $1/\mathsf{poly}(\lambda)$. We show that we can use $\mathcal{A}_0$ to break $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$. On input $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, our reduction proceeds as follows:

1. Sample $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ and let $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}_0 \pmod q$.
2. Run $\mathcal{A}_0$ on input $(\mathbf{A}, \mathbf{y}, |\mathbf{x}_0\rangle\langle\mathbf{x}_0|)$ to obtain a vector $\pi \in \mathbb{Z}_q^n$.
3. Output the vector $(\mathbf{x}_0 - \pi) \in \mathbb{Z}_q^n$.

By assumption, $\mathcal{A}_0$ outputs a valid certificate $\pi \neq \mathbf{x}_0$ with $\|\pi\| \leq \sigma\sqrt{m/2}$ such that $\mathbf{A}\cdot\pi = \mathbf{y} \pmod q$ with probability at least $1/\mathsf{poly}(\lambda)$. Because $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ is sampled from the truncated discrete Gaussian, we have $\|\mathbf{x}_0\| \leq \sigma\sqrt{m/2}$, and thus it holds that

$$\mathbf{A} \cdot (\mathbf{x}_0 - \pi) = \mathbf{0} \pmod q \quad \text{with} \quad \|\mathbf{x}_0 - \pi\| \leq \sigma\sqrt{2m}.$$

Hence, our reduction succeeds at breaking $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ with probability at least $1/\mathsf{poly}(\lambda)$. □

Claim 2. *The probability that the challenger accepts the certificate π in Step 4 of $\mathsf{Hyb}_2(b)$ and the subsequent projective measurement on system C fails (returns the second outcome) is negligible.*

Proof. This follows directly from Claim 1, which implies that except with negligible probability, the register C is in the state

$$|\psi_\pi^{\mathbf{z}}\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \omega_q^{\langle\pi,\mathbf{z}\rangle}|1\rangle)$$

at the time the challenger applies the projective measurement.

□

For any experiment $\mathsf{Hyb}_i(b)$, we define the advantage

$$\mathsf{Adv}(\mathsf{Hyb}_i) := |\Pr[\mathsf{Hyb}_i(0) = 1] - \Pr[\mathsf{Hyb}_i(1) = 1]|.$$

Claim 3.

$$\mathsf{Adv}(\mathsf{Hyb}_2) = 0.$$

Proof. We now consider the following case analysis. First, note that in the case that the challenger rejects because either the certificate is invalid or their projection fails, the experiment does not involve b, and thus the advantage of the adversary is 0. Second, in the case that the challenger's projection succeeds, the register C is either in the state

$$\frac{1}{\sqrt{2}}(|0\rangle + \omega_q^{\langle\pi,\mathbf{z}\rangle}|1\rangle) \quad \text{or} \quad \frac{1}{\sqrt{2}}(|0\rangle - \omega_q^{\langle\pi,\mathbf{z}\rangle}|1\rangle)$$

for some $z \xleftarrow{\$} \mathbb{Z}_q^m$, and thereby completely unentangled from the rest of the system. Notice that the challenger's measurement of system C with outcome c' results in a uniformly random bit, which completely masks b. Therefore, the experiment is also independent of b in this case, and thus the adversary's overall advantage in Hyb_2 is 0. □

Next, we invoke the Gaussian-collapsing property from Theorem 8 to argue the following.

Claim 4.

$$|\mathsf{Adv}(\mathsf{Exp}_2) - \mathsf{Adv}(\mathsf{Exp}_1)| \leq \mathsf{negl}(\lambda).$$

Proof. Recall that Claim 2 shows that the projective measurement performed by the challenger in Step 4 of Hyb_2' succeeds with overwhelming probability. We now argue that the same is also true in Hyb_2. Suppose for the sake of contradiction that there is a non-negligible difference between the success probabilities of the measurement. We now show that this implies the existence of an efficient distinguisher $\mathcal{A}'$ that breaks the Gaussian-collapsing property of the Ajtai hash function family $\mathcal{H} = \{H_\lambda\}_{\lambda \in \mathbb{N}}$ which we showed in Theorem 8. Our reduction proceeds as follows:

$\mathcal{A}'$ receives $(\mathbf{A}, \mathbf{y})$ and a state on register X from its challenger. Next, it samples a random string $z \xleftarrow{\$} \mathbb{Z}_q^m$, prepares a $|+\rangle$ state in system C, applies a controlled-$\mathbf{Z}_q^z$ operation from C to X. Then, it runs $\mathcal{A}_0$ on $(\mathbf{A}, \mathbf{y}, X)$, which outputs a certificate π. Finally, $\mathcal{A}'$ applies the following projective measurement to system C:

$$\left\{ |\psi_\pi^{\mathbf{z}}\rangle\langle\psi_\pi^{\mathbf{z}}|, \mathbb{1} - |\psi_\pi^{\mathbf{z}}\rangle\langle\psi_\pi^{\mathbf{z}}| \right\} \quad \text{with} \quad |\psi_\pi^{\mathbf{z}}\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \omega_q^{\langle \pi, \mathbf{z}\rangle} |1\rangle).$$

and outputs 1 if the measurement succeeds and 0 otherwise. If there is a non-negligible difference in success probabilities of this measurement between $\mathsf{Hyb}_2'(b)$ and $\mathsf{Hyb}_2(b)$ (for any $b \in \{0, 1\}$), then $\mathcal{A}'$ breaks the Gaussian-collapsing property of the Ajtai hash function.

Now, recall that $\mathsf{Hyb}_2(b)$ is identical to $\mathsf{Hyb}_1(b)$, except that the challenger applies an additional a measurement in Step 4. Because the measurement succeeds with overwhelming probability, it follows from the "*Almost As Good As New Lemma*" (Lemma 1) that the advantage of the adversary must remain the same up to a negligible amount. This proves the claim. □

Claim 5.

$$\mathsf{Adv}(\mathsf{Hyb}_1) = \mathsf{Adv}(\mathsf{Hyb}_0)/2.$$

Proof. First note that in $\mathsf{Hyb}_1(b)$, we can imagine measuring register C to obtain c' and aborting if $c' \neq b$ before the challenger sends any information to the adversary. This follows because register C is disjoint from the adversary's registers. Next, by the Pauli-Z twirl property in Lemma 7, we have the following guarantees about the state on system X given to the adversary in $\mathsf{Hyb}_1(b)$.

- In the case $c' = b = 0$, the reduced state on register X is $|\psi_{\mathbf{y}}\rangle$.
- In the case that $c' = b = 1$, the reduced state on register X is a mixture over $\mathbf{x}_0 \sim D_{\mathbb{Z}_q^m, \frac{\sigma}{\sqrt{2}}}$ with $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}_0 \pmod q$, where $\mathbf{x}_0$ results from measuring X in the computational basis.

Thus, this experiment is identical to $\mathsf{Hyb}_0(b)$, except that we decide to abort and output a uniformly random bit b' with probability 1/2 at the beginning of the experiment. □

Putting everything together, we have that $\mathsf{Adv}(\mathsf{Hyb}_0) \leq \mathsf{negl}(\lambda)$, which completes the proof. □

4.4 Public-Key Encryption with Publicly-Verifiable Deletion

In this section, we formalize the cryptographic notion of public-key encryption with publicly-verifiable deletion (PVD). Let us first introduce some relevant definitions.

Definition

We consider public-key encryption schemes with certified deletion for which verification of a deletion certificate is *public*; meaning anyone with access to the verification key can verify that deletion has taken place. The syntax is as follows.

Definition 31 (Public-key encryption with publicly-verifiable deletion)**.** *A public-key encryption scheme with publicly-verifiable deletion (*$\mathsf{PKE}_{\mathsf{PVD}}$*)* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$ *with plaintext space* $\mathcal{M}$ *consists of the following tuple of* QPT *algorithms:*

$\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: *takes as input* 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 (public) verification key* vk *together with a quantum ciphertext* CT*.*

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

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

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

Definition 32 (Correctness of $\mathsf{PKE}_{\mathsf{PVD}}$)**.** *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* $m \in \mathcal{M}$*:*

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

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

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

(Everlasting) certified deletion security

In terms of security, we adopt the following definition introduced in [23, 24].

Definition 33 ((Everlasting certified deletion security)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter and let $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$ be a $\mathsf{PKE}_{\mathsf{PVD}}$ scheme. Let $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ be an adversary. We define the security experiment $\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(b)$ 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, and initializes the algorithm $\mathcal{A}_1$ with its internal state.*
5. *The challenger computes $\mathsf{Vrfy}(\mathsf{vk}, \pi)$. If the output is $\top$, $\mathcal{A}_1$ is run until it outputs b' which is also the output of the experiment. Otherwise, the challenger aborts and $\mathcal{A}$ loses.*

We say that the scheme Σ is $\mathsf{EV\text{-}CD}$-secure if, for any adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ consisting of a QPT algorithm $\mathcal{A}_0$ and a computationally unbounded algorithm $\mathcal{A}_1$, it holds that

$$\mathsf{Adv}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathsf{negl}(\lambda).$$

4.5 Dual-Regev Public-Key Encryption with Publicly-Verifiable Deletion

In this section, we consider the Dual-Regev PKE scheme due to Gentry, Peikert and Vaikuntanathan [68]. Unlike Regev's original PKE scheme in [112], 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.

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 5 (Dual-Regev PKE with Publicly-Verifiable Deletion)**.** *Let $\lambda \in \mathbb{N}$. The Dual-Regev* PKE *scheme* $\mathsf{DualPKE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$ *with PVD is defined as follows:*

- $\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: *sample* $\bar{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ *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* $\mathsf{sk} = (-\bar{\mathbf{x}}, 1)^\intercal \in \mathbb{Z}_q^{m+1}$.

- $\mathsf{Enc}(\mathsf{pk}, x) \to (\mathsf{vk}, |\mathsf{CT}\rangle)$*: 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*

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

  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* $\mathbf{c}^\intercal \cdot \mathsf{sk} \in \mathbb{Z}_q$ *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 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 of decryption and verification of the scheme $\mathsf{DualPKE}_{\mathsf{PVD}}$ in Construction 5.

Lemma 21 (Correctness of decryption)**.** *Let $n \in \mathbb{N}$ and $q \geq 2$ be a prime with $m \geq 2n \log q$, each parameterized by $\lambda \in \mathbb{N}$. 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* $\mathsf{DualPKE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$ *in Construction 5 satisfies:*

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

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 8 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 2 that the ciphertext $|\mathsf{CT}\rangle$ is within negligible trace distance of the state

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

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

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

where $\hat{\mathbf{s}} \xleftarrow{\$} \mathbb{Z}_q^n$ and where $\hat{\mathbf{e}} \sim D_{\mathbb{Z}_q^{m+1}, \frac{\alpha q}{\sqrt{2}}}$ is a sample from the (truncated) discrete Gaussian such that $\|\hat{\mathbf{e}}\| \leq \alpha q \sqrt{\frac{m+1}{2}} < \lfloor \frac{q}{4} \rfloor$. Since $\mathsf{Dec}(\mathsf{sk}, |\mathsf{CT}\rangle)$ computes $\mathbf{c}^\intercal \cdot \mathsf{sk} \in \mathbb{Z} \cap (-\frac{q}{2}, \frac{q}{2}]$ 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 22 (Correctness of verification). *Let $n \in \mathbb{N}$ and $q \geq 2$ be prime with $m \geq 2n \log q$, each parameterized by $\lambda \in \mathbb{N}$. 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* $\mathsf{DualPKE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$ *in Construction 5 satisfies:*

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

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 8 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 2 implies that the Fourier transform of the ciphertext $|\mathsf{CT}\rangle$ is within negligible trace distance of the dual state

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

From Lemma 12, 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 $\|\pi\| \leq \sqrt{m+1}/\sqrt{2}\alpha$. □

Proof of security

Let us now prove the everlasting certified deletion security of our Dual-Regev PKE scheme with PVD in Construction 5.

$\mathsf{IND\text{-}CPA}$ security of $\mathsf{DualPKE}_{\mathsf{PVD}}$. We first prove that our public-key encryption scheme $\mathsf{DualPKE}_{\mathsf{PVD}}$ in Construction 5 satisfies the notion $\mathsf{IND\text{-}CPA}$ security according to Definition 11. The proof follows from Theorem 9 and assumes the hardness of (decisional) LWE (Definition 14). We add it for completeness.

Theorem 11. *Let $n \in \mathbb{N}$ and $q \geq 2$ be prime with $m \geq 2n \log q$, each parameterized by $\lambda \in \mathbb{N}$. 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* $\mathsf{DualPKE}_{\mathsf{PVD}}$ *in Construction 5 is* $\mathsf{IND\text{-}CPA}$*-secure assuming the quantum hardness of the decisional* $\mathsf{LWE}^{m+1}_{n,q,\beta q}$ *problem, for any $\beta \in (0,1)$ with $\alpha/\beta = 2^{o(n)}$.*

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

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

Consider the experiment $\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\Sigma,\mathcal{A},\lambda}(b)$ between $\mathcal{A}$ 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 $\mathsf{Enc}(\mathsf{pk}, m_b)$ 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

$$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{\mathbf{e} \in \mathbb{Z}_q^{m+1}} \varrho_{\alpha q}(\mathbf{e})\, \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle}\, |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal + (0, \dots, 0, m_b \cdot \lfloor q/2 \rfloor) \pmod q\rangle\,. \tag{4.14}$$

Let $\beta \in (0,1)$ be such that $\alpha/\beta = 2^{o(n)}$. From Theorem 9 it follows that, under the $\mathsf{LWE}^m_{n,q,\beta q}$ assumption, the quantum ciphertext $|\mathsf{CT}_b\rangle$ is computationally indistinguishable from the state

$$\sum_{\mathbf{u} \in \mathbb{Z}_q^{m+1}} \omega_q^{-\langle \mathbf{u}, \mathbf{x}_0 \rangle}\, |\mathbf{u}\rangle\,, \qquad \mathbf{x}_0 \sim D_{\Lambda_q^{\mathbf{y}}(\mathbf{A}), \frac{1}{\sqrt{2}\alpha}}. \tag{4.15}$$

Because the state in Eq. (4.17) is completely independent of $b \in \{0,1\}$, it follows that

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

This proves the claim. □

EV-CD security of $\mathsf{DualPKE}_{\mathsf{PVD}}$. In this section, we prove that our public-key encryption scheme $\mathsf{DualPKE}_{\mathsf{PVD}}$ in Construction 5 satisfies the notion of *everlasting certified deletion security* assuming the quantum hardness of LWE and SIS. Our proof makes use of the strengthening of the Gaussian-collapsing property which we proved in Theorem 10.

Theorem 12. *Let $n \in \mathbb{N}$ and $q \geq 2$ be a prime with $m \geq 2n \log q$, each parameterized by $\lambda \in \mathbb{N}$. Let α be a ratio with $\sqrt{8(m+1)} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)}}$. Then,* $\mathsf{DualPKE}_{\mathsf{PVD}}$ *in Construction 5 is* EV-CD*-secure assuming the hardness of* $\mathsf{LWE}_{n,q,\beta q}^{m+1}$ *and* $\mathsf{SIS}_{n,q,\sqrt{2m}/\alpha}^{m+1}$*, for any $\beta \in (0,1)$ with $\alpha/\beta = 2^{o(n)}$.*

Proof. Let $\Sigma = \mathsf{DualPKE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$. We need to show that, for any $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ consisting of a QPT algorithm $\mathcal{A}_0$ and an unbounded algorithm $\mathcal{A}_1$, it holds that

$$\mathsf{Adv}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathsf{negl}(\lambda).$$

Without loss of generality, it suffices to show that the scheme $\hat{\Sigma} = (\mathsf{KeyGen}, \hat{\mathsf{Enc}}, \mathsf{Dec}, \mathsf{Del}, \mathsf{Vrfy})$ is EV-CD-secure. Here, $\hat{\mathsf{Enc}}$ is the same as Enc, except that it additionally applies the Fourier transform to the ciphertext which is output by Enc. We consider the following sequence of hybrids:

$\mathbf{H_0}$: This is the experiment $\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(0)$ 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 $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q]$. Then, the challenger assigns $\mathsf{sk} \leftarrow (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$ as the secret key and $\mathsf{pk} \leftarrow \mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}$ as the public key.
2. $\mathcal{A}_0$ 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
$$\left(\mathsf{vk} \leftarrow (\mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}, \mathbf{y} \in \mathbb{Z}_q^n), \quad |\mathsf{CT}_0\rangle \leftarrow \mathsf{FT}_q |\psi_{\mathbf{y}}\rangle\right).$$
4. At some point in time, $\mathcal{A}_0$ returns a certificate π to the challenger, and initializes $\mathcal{A}_1$ with its internal state.
5. The challenger verifies π and outputs $\top$, if $\mathbf{A} \cdot \pi = \mathbf{y} \pmod q$ and $\|\pi\| \leq \sqrt{m+1}/\sqrt{2}\alpha$. If the output is $\top$, $\mathcal{A}_1$ is run until it outputs b' which is also the output of the experiment. Otherwise, the challenger aborts and $\mathcal{A}$ loses.

$\mathbf{H_1}$: This is same experiment as in $\mathbf{H_0}$, except that (in Step 3) the challenger also measures the ciphertext $|\mathsf{CT}_0\rangle$ in the computational basis before it is send to $\mathcal{A}_0$.

$\mathbf{H_2}$: This is the experiment $\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(1)$ 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 $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q]$. Then, the challenger assigns $\mathsf{sk} \leftarrow (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$ as the secret key and $\mathsf{pk} \leftarrow \mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}$ as the public key.

2. $\mathcal{A}_0$ 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
$$\left(\mathsf{vk} \leftarrow (\mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}, \mathbf{y} \in \mathbb{Z}_q^n), \quad |\mathsf{CT}_1\rangle \leftarrow \mathsf{FT}_q \mathbf{X}_q^{(0,\dots,0,\lfloor \frac{q}{2} \rfloor)} |\psi_{\mathbf{y}}\rangle\right).$$
4. At some point in time, $\mathcal{A}_0$ returns a certificate π to the challenger, and initializes $\mathcal{A}_1$ with its internal state.
5. The challenger verifies π and outputs $\top$, if $\mathbf{A} \cdot \pi = \mathbf{y} \pmod q$ and $\|\pi\| \leq \sqrt{m+1}/\sqrt{2}\alpha$. If the output is $\top$, $\mathcal{A}_1$ is run until it outputs b' which is also the output of the experiment. Otherwise, the challenger aborts and $\mathcal{A}$ loses.

We now show that the hybrids are indistinguishable.

Claim 6.
$$|\Pr[\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(0) = 1] - \Pr[\mathbf{H_1} = 1]| \leq \mathsf{negl}(\lambda).$$

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 8 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 2 implies that the Fourier transform of $\mathsf{FT}_q |\psi_{\mathbf{y}}\rangle$ with $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenPrimal}(\mathbf{A}, 1/\alpha)$ is within negligible trace distance of the dual state
$$|\hat{\psi}_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^{m+1}: \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_{1/\alpha}(\mathbf{x}) |\mathbf{x}\rangle .$$
Therefore, the claim follows immediately from the (everlasting) strong Gaussian-collapsing property in Theorem 10, which implies that the measurement in the computational basis is undetectable. □

Next, we show the following.

Claim 7.
$$|\Pr[\mathbf{H_1} = 1] - \Pr[\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(1) = 1]| \leq \mathsf{negl}(\lambda).$$

Proof. First, recall from Lemma 6 that $\mathsf{FT}_q \mathbf{X}_q^{\mathbf{v}} = \mathbf{Z}_q^{\mathbf{v}} \mathsf{FT}_q$, for all $\mathbf{v} \in \mathbb{Z}_q^m$. Therefore, the proof is the same as in Claim 6, except that the ciphertext output by the challenger in $\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(1)$ is within negligible trace distance of the quantum state
$$\mathbf{Z}_q^{\langle \mathbf{x},(0,\dots,0,\lfloor \frac{q}{2} \rfloor)\rangle} |\hat{\psi}_{\mathbf{y}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^{m+1}: \\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \varrho_{1/\alpha}(\mathbf{x}) \omega_q^{\langle \mathbf{x},(0,\dots,0,\lfloor \frac{q}{2} \rfloor)\rangle} |\mathbf{x}\rangle .$$

Therefore, the claim follows from the (everlasting) strong Gaussian-collapsing property in Theorem 10, since applying the phase operator $\mathbf{Z}_q^{(0,\dots,0,\lfloor \frac{q}{2} \rfloor)}$ before the measurement does not affect the measurement outcome in the computational basis. □

Because the hybrids $\mathbf{H_0}$ and $\mathbf{H_2}$ are indistinguishable, this implies that

$$\mathsf{Adv}_{\hat{\Sigma},\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda).$$

This proves the claim. □

Next, we show how to extend our Dual-Regev PKE scheme with certified deletion in Construction 5 to a fully homomorphic encryption scheme of the same type.

4.6 Fully Homomorphic Encryption with Publicly-Verifiable Deletion

In this section, we formalize the notion of homomorphic encryption with publicly-verifiable deletion which enables an untrusted quantum server to compute on encrypted data and to also prove data deletion to a client. Let us first introduce some relevant definitions.

Definition

We begin with the following definition.

Definition 34 (Homomorphic encryption with publicly-verifiable deletion). *Let $\lambda \in \mathbb{N}$ be the security parameter. A homomorphic encryption scheme with publicly-verifiable deletion is a tuple* $\mathsf{HE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *of* QPT *algorithms:*

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

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

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

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

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

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

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 35 (Compactness and full homomorphism)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. A homomorphic encryption scheme* $\mathsf{HE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *with publicly-verifiable deletion 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*

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

We say that a fully homomorphic encryption scheme with certified deletion $(\mathsf{FHE}_{\mathsf{PVD}})$ *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 36 (Correctness of verification)**.** *A homomorphic encryption scheme with certified deletion* $\mathsf{HE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *has correctness of verification if the following property holds for any integer $\lambda \in \mathbb{N}$ and any set of inputs $x = (x_1, \ldots, x_\ell) \in \{0,1\}^\ell$*

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

(Everlasting) certified deletion security

Our notion of certified deletion security for homomorphic encryption (HE) schemes is identical to the notion of EV-CD security for public-key encryption schemes in Definition 33.

Definition 37 ((Everlasting certified deletion security)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter and let* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *be a homomorphic encryption scheme with publicly-verifiable deletion. Let $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ be an adversary. We define the security experiment* $\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(b)$ *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, and initializes the algorithm $\mathcal{A}_1$ with its internal state.*
5. *The challenger computes* $\mathsf{Vrfy}(\mathsf{vk}, \pi)$. *If the output is $\top$, $\mathcal{A}_1$ is run until it outputs b' which is also the output of the experiment. Otherwise, the challenger aborts and $\mathcal{A}$ loses.*

We say that the scheme Σ is EV-CD*-secure if, for any adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ consisting of a* QPT *algorithm $\mathcal{A}_0$ and a computationally unbounded algorithm $\mathcal{A}_1$, it holds that*

$$\mathsf{Adv}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(0) = 1] - \Pr[\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(1) = 1]| \leq \mathsf{negl}(\lambda).$$

4.7 Dual-Regev Fully Homomorphic Encryption with Publicly-Verifiable Deletion

In this section, we describe a protocol that allows an untrusted quantum server to perform homomorphic operations on encrypted data and to also 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 4.4) – 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 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 4.5, we give an example of a Boolean circuit composed of three NAND gates that takes as input a string $x \in \{0, 1\}^4$.

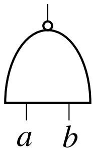

Figure 4.4: NAND gate.

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 [69] in the context of homomorphic encryption and is also featured in the Dual-Regev FHE scheme of Mahadev [99]. Let $n \in \mathbb{N}$, $q \geq 2$ be a prime modulus and $m \geq 2n \log q$. Then, for $N = (m + 1) \cdot \lceil \log q \rceil$, we let $\mathbb{1}$ be the $(m + 1) \times (m + 1)$ identity matrix and

$$\mathbf{G} = [\mathbb{1} \,\|\, 2\mathbb{1} \,\|\, \ldots \,\|\, 2^{\lceil \log q \rceil - 1}\mathbb{1}] \in \mathbb{Z}_q^{(m+1)\times N}$$

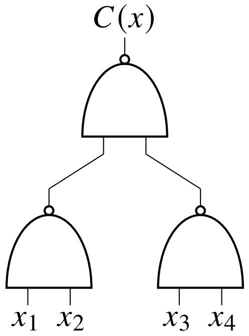

Figure 4.5: A Boolean circuit C made up of three NAND gates which takes as input a binary string of the form $x \in \{0,1\}^4$. The top-most NAND gate is the designated output node with outcome $C(x) \in \{0,1\}$.

denote the so-called *gadget matrix* which converts a binary representation of a vector back to its original vector representation over the field $\mathbb{Z}_q$. Note that the associated (non-linear) inverse operation $\mathbf{G}^{-1}$ converts vectors in $\mathbb{Z}_q^{m+1}$ to their binary representation in $\{0,1\}^N$. In other words, we have that $\mathbf{G} \circ \mathbf{G}^{-1}$ acts as the identity operator.

Our (leveled) FHE scheme with certified deletion is based on the (leveled) Dual-Regev FHE scheme introduced by Mahadev [99] which is a variant of the LWE-based FHE scheme proposed by Gentry, Sahai and Waters [69]. We base our choice of parameters on the aforementioned two works.

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

Construction 6 (Dual-Regev leveled FHE). *Let $\lambda \in \mathbb{N}$ be the security parameter and let L be an upper bound on the* NAND *depth of the circuit which is to be evaluated. The Dual-Regev leveled* FHE *scheme* DualFHE = (KeyGen, Enc, Dec, Eval) *consists of the following* PPT *algorithms:*

KeyGen$(1^\lambda, 1^L) \rightarrow$ (pk, sk) : *sample* $\bar{\mathbf{A}} \overset{\$}{\leftarrow} \mathbb{Z}_q^{n \times m}$ *and* $\bar{\mathbf{x}} \overset{\$}{\leftarrow} \{0,1\}^m$ *and let* $\mathbf{A} = [\bar{\mathbf{A}} | \bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q]^\intercal$. *Output* (pk, sk), *where* pk $= \mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}$ *and* sk $= (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$.

Enc(pk, x) : *to encrypt* $x \in \{0,1\}$, *parse* $\mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n} \leftarrow$ pk, *sample* $\mathbf{S} \overset{\$}{\leftarrow} \mathbb{Z}_q^{n \times N}$ *and an error* $\mathbf{E} \sim D_{\mathbb{Z}^{(m+1)\times N}, \alpha q}$ *and output* 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 power-of-two gadget matrix in Eq.* (4.7).

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

- *parse the ciphertext tuple as* (CT$_1, \ldots,$ CT$_\ell$) $\leftarrow$ CT.
- *repeat for every* NAND *gate in C: to apply a* NAND *gate on a ciphertext pair* (CT$_i$, CT$_j$), *parse matrices* $\mathbf{C}_i \leftarrow$ CT$_i$ *and* $\mathbf{C}_j \leftarrow$ 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* $c = \mathsf{sk}^\intercal \cdot \mathbf{c}_N \in \mathbb{Z} \cap (-\frac{q}{2}, \frac{q}{2}]$, *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 [99, Theorem 5.1].

Theorem 13 ([99])**.** *Let* $\lambda \in \mathbb{N}$ *be the security parameter. Let* $n \in \mathbb{N}$, *let* $q \geq 2$ *be a prime modulus and* $m \geq 2n \log q$, *each parameterized by* λ. *Let* $N = (m+1) \cdot \lceil \log q \rceil$ *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*

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

Then, the scheme in Construction 6 is an IND-CPA*-secure leveled fully homomorphic encryption scheme under the* $\mathsf{LWE}_{n,q,\alpha q}^{m+1}$ *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 [37].

The leveled Dual-Regev FHE scheme inherits a crucial property from its public-key counterpart. Namely, in contrast to the FHE scheme in [69], 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 6, 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 $\|\mathbf{E}\| \leq \alpha q \sqrt{(m+1)} \cdot (N+1)^L$ (see [69] for details). Choosing $1/\alpha$ to be sub-exponential in N as in [69], 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 $\|\mathbf{E}\| \leq \frac{q}{4\sqrt{(m+1)}}$.

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 $\mathsf{FHE}_{\mathsf{PVD}}$. 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 [99] which is a variant of the scheme due to Gentry, Sahai and Waters [69].

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 λ).

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

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

Construction 7 (Dual-Regev leveled FHE scheme with publicly-verifiable 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 6. The Dual-Regev (leveled)* FHE *scheme* $\mathsf{DualFHE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *with publicly-verifiable deletion is defined by the following* QPT *algorithms:*

$\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: *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 $(|\psi_{\mathbf{y}_i}\rangle, \mathbf{y}_i) \leftarrow \mathsf{GenPrimal}(\mathbf{A}^\intercal, 1/\alpha)$ in Algorithm 2, where $\mathbf{y}_i \in \mathbb{Z}_q^n$, and output the pair*

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

where $(\mathbf{g}_1, \dots, \mathbf{g}_N)$ are the columns of the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ in Eq. (4.7).

$\mathsf{Eval}(C, |\mathsf{CT}\rangle) \rightarrow |\widetilde{\mathsf{CT}}\rangle$: *apply the Boolean circuit C composed of* NAND *gates to the ciphertext* $|\mathsf{CT}\rangle$ *in system* $C_{\mathsf{in}} = C_1 \cdots C_\ell$ *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 38 to systems* $C_i C_j$ *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}*.*

Output $|\widetilde{\mathsf{CT}}\rangle$*, where* $|\widetilde{\mathsf{CT}}\rangle$ *is the final post-evaluation state in systems* $C_{\mathsf{in}} C_{\mathsf{aux}} C_{\mathsf{out}}$ *and*

- $C_{\mathsf{in}} = C_1 \cdots C_\ell$ *denotes the initial ciphertext systems of* $|\mathsf{CT}_1\rangle \otimes \cdots \otimes |\mathsf{CT}_\ell\rangle$*.*
- C_{aux} *denotes all intermediate auxiliary ciphertext systems.*
- C_{out} *denotes the final ciphertext system corresponding to the output of the circuit C.*

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

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

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

Let us now define how to perform the homomorphic NAND gate in Construction 7 in more detail.

Definition 38 (Homomorphic NAND gate)**.** *Let $q \geq 2$ be a modulus, and let $m, N \in \mathbb{N}$. 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*

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

where $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ *is the gadget matrix in Eq. (4.7).*

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

Example. Consider a pair of ciphertexts $|\mathsf{CT}_i\rangle \otimes |\mathsf{CT}_j\rangle$ which encrypt two bits $x_i, x_j \in \{0,1\}$ as in Construction 7. Let $U_{\mathsf{NAND}_{ij}}$ denote the homomorphic NAND gate applied to systems $C_i C_j$. Then,

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

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

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

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 7. 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,

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

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 $\|\mathbf{E}_i\|, \|\mathbf{E}_j\| \leq \alpha q\sqrt{(m+1)/2}$, it follows that the resulting error after a single NAND gate is at most (see also [69, 99] for more details)

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

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 [99].

Proof of correctness. Let us now verify the correctness of decryption and verification of Construction 7.

Lemma 23 (Compactness and full homomorphism of $\mathsf{DualFHE}_{\mathsf{PVD}}$)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n \in \mathbb{N}$, let $q \geq 2$ be a prime and $m \geq 2n \log q$, each parameterized by λ. Let $N = (m+1) \cdot \lceil \log q \rceil$ 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*

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

Then, the scheme $\mathsf{DualFHE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *in Construction 7 is a compact and fully homomorphic encryption scheme with publicly-verifiable deletion. In other words, for any efficienty computable circuit $C : \{0,1\}^\ell \rightarrow \{0,1\}$ and any set of inputs $x = (x_1, \ldots, x_\ell) \in \{0,1\}^\ell$, it holds that:*

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

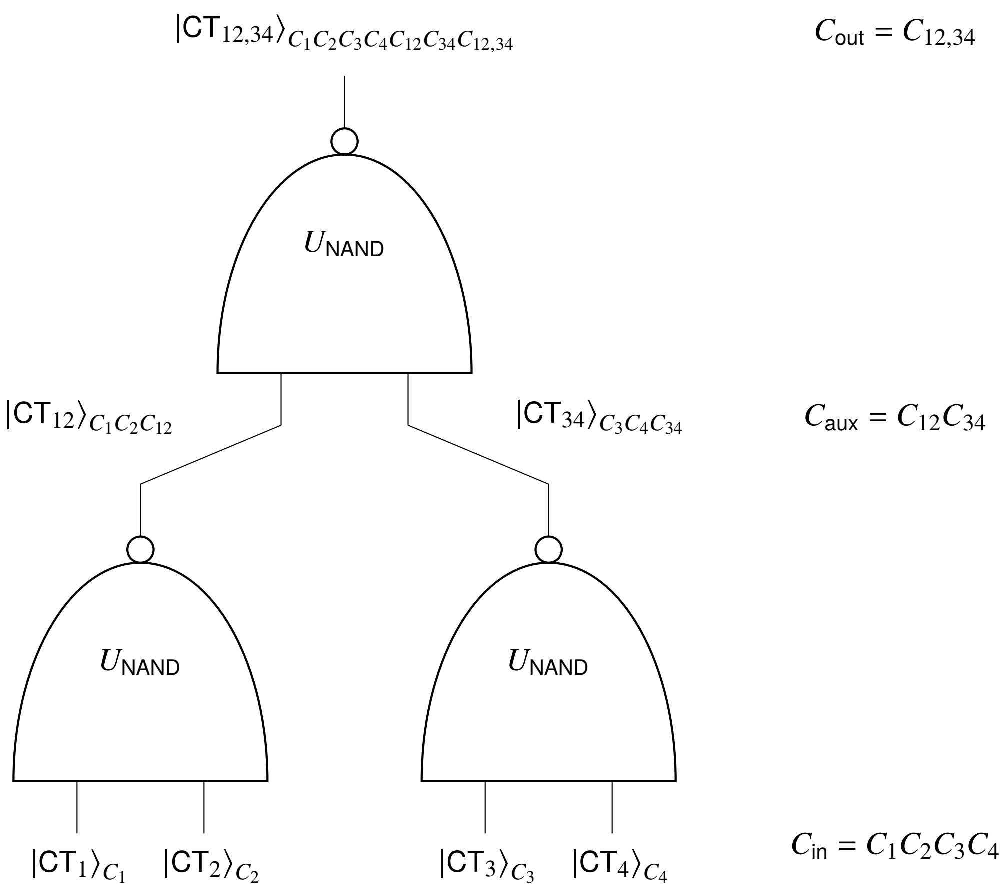
$|\mathsf{CT}_{12,34}\rangle_{C_1C_2C_3C_4C_{12}C_{34}C_{12,34}}$
$C_{\mathsf{out}} = C_{12,34}$
U_{NAND}
$|\mathsf{CT}_{12}\rangle_{C_1C_2C_{12}}$
$|\mathsf{CT}_{34}\rangle_{C_3C_4C_{34}}$
$C_{\mathsf{aux}} = C_{12}C_{34}$
U_{NAND}
U_{NAND}
$|\mathsf{CT}_1\rangle_{C_1}$
$|\mathsf{CT}_2\rangle_{C_2}$
$|\mathsf{CT}_3\rangle_{C_3}$
$|\mathsf{CT}_4\rangle_{C_4}$
$C_{\mathsf{in}} = C_1C_2C_3C_4$

Figure 4.6: 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 7. 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 $C(x) \in \{0,1\}$.

Proof. Let $|\mathsf{CT}\rangle$ be the ciphertext output by $\mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Enc}(\mathsf{pk}, x)$, where $x \in \{0,1\}^\ell$ denotes the plaintext, and let $(|\widetilde{\mathsf{CT}}\rangle, t_C) \leftarrow \mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Eval}(C, |\mathsf{CT}\rangle)$ 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 $\{\mathbf{X} \in \mathbb{Z}_q^{(m+1)\times N} : \|\mathbf{X}\| \leq \alpha q \sqrt{(m+1)/2}\}$. Recall that $\mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Dec}(\mathsf{sk}, |\widetilde{\mathsf{CT}}\rangle)$ 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

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

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 38 is at most $\alpha q\sqrt{(m+1)/2}(N+1)^L$, which is less than $\frac{q}{4\sqrt{(m+1)}}$ by our choice of parameters. Therefore, the procedure $\mathsf{DualFHE.Dec}_{\mathsf{sk}}$ decrypts a computational basis state in system C_{out} of the state $|\widetilde{\mathsf{CT}}\rangle$ correctly with probability at least $1-\mathsf{negl}(\lambda)$. Furthermore, because the procedure $\mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Dec}$ is independent of the circuit C and its depth L, the scheme $\mathsf{DualFHE}_{\mathsf{PVD}}$ is compact. This proves the claim. □

Let us now verify the correctness of verification of the scheme $\mathsf{DualFHE}_{\mathsf{PVD}}$ in Construction 7 according to Definition 36. We show the following.

Lemma 24 (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 $N = (m+1)\cdot\lceil\log q\rceil$ 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*

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

Then, the Dual-Regev FHE *scheme* $\mathsf{DualFHE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *with certified deletion in Construction 7 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:*

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

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 8 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 $(\mathbf{g}_1,\dots,\mathbf{g}_N)$ are the columns of the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times N}$ in Eq. (4.7). Given our choice,

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

Corollary 2 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 12, it follows that the distribution of computational basis measurement outcomes is within negligible total variation distance of the sample

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

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

Proof of security

Let us now analyze the security of our FHE scheme with publicly-verifiable deletion in Construction 7. Note that the results in this section all essentially carry over from Section 4.5, where we analyzed the security of our Dual-Regev PKE scheme with publicly-verifiable deletion.

$\mathsf{IND\text{-}CPA}$ **security of** $\mathsf{DualFHE}_{\mathsf{PVD}}$**.** We first prove that our scheme $\mathsf{FHE}_{\mathsf{PVD}}$ in Construction 7 satisfies the notion $\mathsf{IND\text{-}CPA}$ security according to Definition 11. The proof is identical to the proof of $\mathsf{IND\text{-}CPA}$-security of our $\mathsf{DualPKE}$ scheme in Theorem 11. We add it for completeness.

Theorem 14. *Let $n \in \mathbb{N}$, let $q \geq 2$ be a modulus, let $m \geq 2n \log q$ and let $N = (m+1)\lceil \log q \rceil$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\alpha \in (0,1)$ be a noise ratio parameter such that $\sqrt{8(m+1)} \leq \frac{1}{\alpha} \leq \frac{q}{\sqrt{8(m+1)}}$. Then, the scheme $\mathsf{DualFHE}_{\mathsf{PVD}}$ in Construction 7 is $\mathsf{IND\text{-}CPA}$-secure assuming the quantum hardness of $\mathsf{LWE}^{m+1}_{n,q,\beta q}$, for any $\beta \in (0,1)$ with $\alpha/\beta = 2^{o(n)}$.*

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

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

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 $(\mathsf{vk}, \mathsf{CT}_b) \leftarrow \mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Enc}(\mathsf{pk}, m_b)$, 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

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

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

$$\sum_{\mathbf{S}\in\mathbb{Z}_q^{n\times N}} \sum_{\mathbf{E}\in\mathbb{Z}_q^{(m+1)\times N}} \varrho_{\alpha q}(\mathbf{E})\, \omega_q^{-\mathrm{Tr}[\mathbf{S}^T\mathbf{Y}]}\, |\mathbf{A}\cdot\mathbf{S}+\mathbf{E}+m_b\cdot\mathbf{G} \;(\mathrm{mod}\; q)\rangle\,. \tag{4.16}$$

Here, $\mathbf{Y} \in \mathbb{Z}_q^{n\times N}$ is the matrix composed of the columns $\mathbf{y}_1,\dots,\mathbf{y}_N$. Let $\beta \in (0,1)$ be any parameter with $\alpha/\beta = 2^{o(n)}$. Then, it follows from Theorem 9 that, under the (decisional) $\mathsf{LWE}_{n,q,\beta q}^{m+1}$ assumption, $|\mathsf{CT}_b\rangle$ is computationally indistinguishable from the state

$$\sum_{\mathbf{U}\in\mathbb{Z}_q^{(m+1)\times N}} \omega_q^{\mathrm{Tr}[\mathbf{U}^T\bar{\mathbf{X}}]}\, |\mathbf{U}\rangle\,, \qquad \bar{\mathbf{X}} = (\bar{\mathbf{x}}_1,\dots,\bar{\mathbf{x}}_N) \sim D_{\Lambda_q^{\mathbf{y}_1}(\mathbf{A}),\frac{1}{\sqrt{2}\alpha}} \times\cdots\times D_{\Lambda_q^{\mathbf{y}_N}(\mathbf{A}),\frac{1}{\sqrt{2}\alpha}}\,. \tag{4.17}$$

Here $(\bar{\mathbf{x}}_1,\dots,\bar{\mathbf{x}}_N)$ refer to the columns of the matrix $\bar{\mathbf{X}} \in \mathbb{Z}_q^{(m+1)\times N}$. Finally, because the state in Eq. (4.17) is completely independent of the bit $b \in \{0,1\}$, it follows that

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

This proves the claim. □

EV-CD security of DualFHE$_{\mathsf{PVD}}$. Let us now analyze the security of our Dual-Regev homomorphic encryption scheme $\mathsf{DualFHE}_{\mathsf{PVD}}$ in Construction 7. The proof is similar to the proof of Theorem 12. We add it for completeness.

Theorem 15. *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 $N = (m+1)\cdot\lceil\log q\rceil$ and let L be an upper bound on the depth of the $\mathsf{poly}(\lambda)$-sized Boolean circuit which is to be evaluated. Let $\alpha \in (0,1)$ be a noise ratio such that*

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

Then, the Dual-Regev homomorphic encryption scheme $\mathsf{DualFHE}_{\mathsf{PVD}}$ in Construction 7 is EV-CD-secure assuming the hardness of $\mathsf{LWE}_{n,q,\beta q}^{m+1}$ and $\mathsf{SIS}_{n,q,\sqrt{2m}/\alpha}^{m+1}$, for any $\beta \in (0,1)$ with $\alpha/\beta = 2^{o(n)}$.

Proof. Let $\Sigma = \mathsf{DualFHE}_{\mathsf{PVD}} = (\mathsf{KeyGen},\mathsf{Enc},\mathsf{Dec},\mathsf{Eval},\mathsf{Del},\mathsf{Vrfy})$. We need to show that, for any $\mathcal{A} = (\mathcal{A}_0,\mathcal{A}_1)$ consisting of a QPT algorithm $\mathcal{A}_0$ and an unbounded algorithm $\mathcal{A}_1$, it holds that

$$\mathsf{Adv}_{\Sigma,\mathcal{A}}(\lambda) := |\Pr[\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(0)=1] - \Pr[\mathsf{EvExp}_{\Sigma,\mathcal{A},\lambda}(1)=1]| \leq \mathsf{negl}(\lambda).$$

Without loss of generality, it suffices to show that the scheme $\hat{\Sigma} = (\mathsf{KeyGen},\hat{\mathsf{Enc}},\mathsf{Dec},\mathsf{Del},\mathsf{Vrfy})$ is EV-CD-secure. Here, $\hat{\mathsf{Enc}}$ is the same as Enc, except that it additionally applies the Fourier transform to the ciphertext which is output by Enc. We consider the following sequence of hybrids:

$\mathbf{H_0}$: This is the experiment $\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(0)$ 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 $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q)]^\intercal$. The challenger then assigns $\mathsf{sk} \leftarrow (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$ as the secret key and $\mathsf{pk} \leftarrow \mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}$ as the public key.
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 generates a sequence of pairs $(|\psi_{\mathbf{y}_i}\rangle, \mathbf{y}_i) \leftarrow \mathsf{GenPrimal}(\mathbf{A}^\intercal, 1/\alpha)$ in Algorithm 2, for each $i \in [N]$, and sends the following to the adversary $\mathcal{A}_0$:

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

4. At some point in time, $\mathcal{A}$ returns a certificate $\pi = (\pi_1, \dots, \pi_N)$ to the challenger, and initializes $\mathcal{A}_1$ with its internal state.
5. The challenger checks if $\mathbf{A}^\intercal \cdot \pi_i = \mathbf{y}_i \pmod q$ and $\|\pi_i\| \leq \sqrt{m+1}/\sqrt{2}\alpha$ for $i \in [N]$. If the output is $\top$ for each $i \in [N]$, $\mathcal{A}_1$ is run until it outputs b' which is also the output of the experiment. Otherwise, the challenger aborts and $\mathcal{A}$ loses.

$\mathbf{H_1}$: This is same experiment as in $\mathbf{H_0}$, except that (in Step 3) the challenger also measures the ciphertext $|\mathsf{CT}_0\rangle$ in the computational basis before it is send to $\mathcal{A}_0$.

$\mathbf{H_2}$: This is the experiment $\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(1)$ 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 $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q)]^\intercal$. The challenger then assigns $\mathsf{sk} \leftarrow (-\bar{\mathbf{x}}, 1) \in \mathbb{Z}_q^{m+1}$ as the secret key and $\mathsf{pk} \leftarrow \mathbf{A} \in \mathbb{Z}_q^{(m+1)\times n}$ as the public key.
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 generates a sequence of pairs $(|\psi_{\mathbf{y}_i}\rangle, \mathbf{y}_i) \leftarrow \mathsf{GenPrimal}(\mathbf{A}^\intercal, 1/\alpha)$ in Algorithm 2, for each $i \in [N]$, and sends the following pair to the adversary $\mathcal{A}_0$:

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

4. At some point in time, $\mathcal{A}$ returns a certificate $\pi = (\pi_1, \dots, \pi_N)$ to the challenger, and initializes $\mathcal{A}_1$ with its internal state.

5. The challenger checks if $\mathbf{A}^\intercal \cdot \pi_i = \mathbf{y}_i \pmod q$ and $\|\pi_i\| \leq \sqrt{m+1}/\sqrt{2}\alpha$ for $i \in [N]$. If the output is $\top$ for each $i \in [N]$, $\mathcal{A}_1$ is run until it outputs b' which is also the output of the experiment. Otherwise, the challenger aborts and $\mathcal{A}$ loses.

We now show that the hybrids are indistinguishable.

Claim 8.

$$|\Pr[\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(0) = 1] - \Pr[\mathbf{H_1} = 1]| \leq \mathsf{negl}(\lambda).$$

Proof. By the Leftover Hash Lemma (Lemma 4), the distribution of $\mathbf{A} = [\bar{\mathbf{A}}|\bar{\mathbf{A}} \cdot \bar{\mathbf{x}} \pmod q]^\intercal$ is within negligible total variation distance of the uniform distribution over $\mathbb{Z}_q^{(m+1)\times n}$. From Lemma 8 it follows that the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$ with overwhelming probability. Since $\alpha \in (0,1)$ is a noise ratio parameter with

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

Corollary 2 implies that the Fourier transform of $(|\psi_{\mathbf{y}_i}\rangle, \mathbf{y}_i) \leftarrow \mathsf{GenPrimal}(\mathbf{A}^\intercal, 1/\alpha)$ in Algorithm 2, for each $i \in [N]$, is within negligible trace distance of the dual state

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

Therefore, the claim follows immediately from the (everlasting) strong Gaussian-collapsing property in Theorem 10, which implies that the measurement in the computational basis is undetectable. □

Next, we show the following.

Claim 9.

$$|\Pr[\mathbf{H_1} = 1] - \Pr[\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(1) = 1]| \leq \mathsf{negl}(\lambda).$$

Proof. First, recall from Lemma 6 that $\mathsf{FT}_q \mathbf{X}_q^{\mathbf{v}} = \mathbf{Z}_q^{\mathbf{v}} \mathsf{FT}_q$, for all $\mathbf{v} \in \mathbb{Z}_q^m$. Therefore, the proof is the same as in Claim 8, except that the ciphertext output by the challenger in $\mathsf{EvExp}_{\hat{\Sigma},\mathcal{A},\lambda}(1)$ is within negligible trace distance of the quantum state

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

Therefore, the claim follows from the (everlasting) strong Gaussian-collapsing property in Theorem 10, since applying the phase operators $\mathbf{Z}_q^{\mathbf{g}_i}$, for $i \in [N]$, before the measurement does not affect the measurement outcome in the computational basis. □

Because the hybrids $\mathbf{H_0}$ and $\mathbf{H_2}$ are indistinguishable, this implies that

$$\mathsf{Adv}_{\hat{\Sigma},\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda).$$

This proves the claim. □

4.8 Four-Message Protocol for FHE with Simultaneous Data Deletion

Our Dual-Regev encryption scheme in Construction 7 separately supports both homomorphic operations as well as publicly-verifiable deletion. It is therefore natural to ask whether it is possible to achieve both tasks simultaneously, say in a protocol between a client and an untrusted server. Remarkably, such a protocol would allow an untrusted server to compute on private data and, if requested, to simultaneously also prove data deletion to a client. We now show that such a protocol is indeed possible, albeit with a few important caveats which we explain below.

Recall that Construction 7 has the property that when applying homomorphic operations, the resulting ciphertext 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. Unfortunately, the resulting ciphertext 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 server perform both homomorphic computations and, if requested, to later 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 prevents us from obtaining a valid deletion certificate via a Fourier basis measurement. Our solution to the problem involves one additional round of interaction (as compared to a regular homomorphic encryption protocol) between the quantum server and the client in order to *certify deletion*.

In Protocol 1, we describe a four-message protocol for FHE with simultaneous data deletion which is based on our Dual-Regev (leveled) fully homomorphic encryption scheme in Construction 7. This is a four-message interactive protocol which enables an untrusted quantum server to compute on encrypted data and, if requested, to simultaneously prove data deletion to a client. The basic idea behind our approach for certified deletion is the following: After performing a Boolean circuit C via a sequence of U_{NAND} gates starting from the ciphertext $|\mathsf{CT}\rangle$ in system C_{in} corresponding to an encryption of $x \in \{0,1\}^{\ell}$, the server simply sends the quantum system C_{out} containing an encryption of $C(x)$ to the client. Then, using the secret key sk (i.e., a trapdoor for the public matrix pk), it is possible for the client to *extract* the outcome $C(x)$ from the system C_{out} with overwhelming probability without significantly damaging the state. In Lemma 25, we show that it is possible to rewind the procedure in a way that it results in a state which is negligibly close to the original state in system C_{out}. At this step of the protocol, the client has learned the outcome of the homomorphic application of the circuit C while the server 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 enable *certified deletion*, the client must now return the system C_{out} to the server. Having access to all three systems $C_{\mathsf{in}}C_{\mathsf{aux}}C_{\mathsf{out}}$, the server 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 server's possession is now approximately a simple product of Gaussian states, the server 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. Crucially, this requires that the server is *honest* during the evaluation phase of the protocol. We further elaborate on this important caveat of our protocol towards the end of this section.

Let us now describe our four-message protocol for FHE with simultaneous data deletion, which is based on our Dual-Regev homomorphic encryption scheme in Construction 7.

Protocol 1 (Four-Message Protocol for FHE with Simultaneous Data Deletion)**.**
Let $\lambda \in \mathbb{N}$ *be the security parameter and let* $\mathsf{DualFHE}_{\mathsf{PVD}} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Del}, \mathsf{Vrfy})$ *be the Dual-Regev (leveled)* FHE *scheme with publicly-verifiable deletion in Construction 7. Consider the following quantum interactive protocol* $\Pi = \langle C(1^\lambda, x), \mathcal{S}(1^\lambda)\rangle$ *between a quantum client* C *with data* $x \in \{0,1\}^\ell$ *and a quantum server* $\mathcal{S}$*:*

Setup phase:

1. C *runs* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{KeyGen}(1^\lambda, 1^L)$ *to generate a pair of keys, where* L *is an upper bound on the depth of the Boolean circuit which is to be evaluated.*

Encryption phase:

1. C *runs the procedure* $(|\mathsf{CT}\rangle, \mathsf{vk}) \leftarrow \mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Enc}(\mathsf{pk}, x)$ *to obtain a quantum ciphertext* $|\mathsf{CT}\rangle$ *and a public verification key* vk.
2. C *sends the public key* pk *and ciphertext* $|\mathsf{CT}\rangle$ *to the quantum server* $\mathcal{S}$.

Evaluation phase:

1. $\mathcal{S}$ *runs the evaluation procedure* $|\widetilde{\mathsf{CT}}\rangle \leftarrow \mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Eval}(C, |\mathsf{CT}\rangle)$*, for some classical Boolean circuit* C*, which results in a quantum ciphertext* $|\widetilde{\mathsf{CT}}\rangle$ *in systems* $C_{\mathsf{in}}C_{\mathsf{aux}}C_{\mathsf{out}}$.
2. $\mathcal{S}$ *sends the register* C_{out} *to* C.

3. *$\mathcal{C}$ appends an ancilla register $|0\rangle_M$ and then runs the procedure* $\mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Dec}(\mathsf{sk},\cdot)$ *coherently on register C_{out}. Then, $\mathcal{C}$ measures system M to obtain a bit $y \in \{0,1\}$ (the supposed output of the Boolean circuit C).*
4. *$\mathcal{C}$ returns the post-measurement register $\widetilde{C_{\mathsf{out}}}$ to $\mathcal{S}$.*
5. *$\mathcal{S}$ applies the previous evaluation procedure* $\mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Eval}(C,\cdot)$ *in reverse on input the registers $C_{\mathsf{in}}C_{\mathsf{aux}}\widetilde{C_{\mathsf{out}}}$. Let $\hat{\mathsf{CT}}$ denote the resulting ciphertext.*

Deletion phase:

- *$\mathcal{S}$ runs the procedure $\pi \leftarrow$* $\mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Del}(\hat{\mathsf{CT}})$.
- *$\mathcal{C}$ runs $\texttt{flag} \leftarrow$* $\mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Vrfy}(\mathsf{vk},\pi)$ *and outputs $\texttt{flag} \in \{\top,\bot\}$.*

Correctness

Recall that the procedure $\mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Eval}$ in Construction 7 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 in order to allow the server to subsequently prove data deletion to a client.

Lemma 25 (Rewinding lemma)**.** *Let $\lambda \in \mathbb{N}$. Let $n \in \mathbb{N}$, let $q \geq 2$ be a prime modulus and $m \geq 2n\log q$. Let $N = (m+1)\cdot\lceil \log q\rceil$ 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*

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

Let $\mathsf{DualFHE}_{\mathsf{PVD}} = (\mathsf{KeyGen},\mathsf{Enc},\mathsf{Dec},\mathsf{Eval},\mathsf{Del},\mathsf{Vrfy})$ *be the Dual-Regev (leveled)* FHE *scheme with publicly-verifiable deletion in Construction 7 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: After the evaluation phase of the protocol $\Pi = \langle \mathcal{C}(1^\lambda,x),\mathcal{S}(1^\lambda)\rangle$ is complete, the server $\mathcal{S}$ is in possession of a quantum state $\hat{\mathsf{CT}}$ in system C_{in} that satisfies*

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

where $|\mathsf{CT}\rangle \leftarrow \mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Enc}(\mathsf{pk},x)$ *is the initial quantum ciphertext in systems $C_{\mathsf{in}}C_{\mathsf{aux}}C_{\mathsf{out}}$ for the pair of keys* $(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{KeyGen}(1^\lambda,1^L)$.

Proof. Let $\lambda \in \mathbb{N}$, $x \in \{0,1\}^\ell$ be a plaintext and C be any Boolean circuit of NAND-depth $L = \mathsf{poly}(\lambda)$. Let $|\widetilde{\mathsf{CT}}\rangle \leftarrow \mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Eval}(C,|\mathsf{CT}\rangle)$ be the post-evaluation state $|\widetilde{\mathsf{CT}}\rangle$ in systems $C_{\mathsf{in}}C_{\mathsf{aux}}C_{\mathsf{out}}$ and let $\hat{\mathsf{CT}}$ be the state at the end of the evaluation phase of the protocol $\Pi = \langle \mathcal{C}(1^\lambda,x),\mathcal{S}(1^\lambda)\rangle$. Recall that, in Lemma 23, we established that there exists a negligible

$\varepsilon(\lambda)$ such that $\mathsf{DualFHE.Dec}_{\mathsf{sk}}$ 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 $U_{\mathsf{DualFHE.Dec}_{\mathsf{sk}}}$, measuring the ancillary register M and rewinding the computation, results in a mixed state ϱ that is within trace distance $\sqrt{\varepsilon}$ of the post-evaluation state $|\widetilde{\mathsf{CT}}\rangle$. Let U_{t_C} be sequence of homomorphic NAND gates which are applied according to the circuit C. Notice that, by reversing the sequence of NAND gates which are applied 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 $\hat{\mathsf{CT}} = U_{t_C}^\dagger \varrho\, U_{t_C}$. Therefore,

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

since the trace distance is unitarily invariant. Thus, $\varepsilon(\lambda) = \mathsf{negl}(\lambda)$ which proves the claim. □

Security. Our four-message protocol for FHE with simultaneous data deletion in Protocol 1 immediately inherits the following security guarantees from the underlying Dual-Regev homorphic encryption scheme Construction 7:

- **Privacy:** This ensures that the data is computationally hidden from the view of the server, once the encryption phase ends. In other words, for any pair of messages $x_0, x_1 \in \{0,1\}^\ell$ (selectively chosen by the server), the following are computationally indistinguishable,

  $$\mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Enc}(\mathsf{pk}, x_0) \approx_c \mathsf{DualFHE}_{\mathsf{PVD}}.\mathsf{Enc}(\mathsf{pk}, x_1),$$

  where pk is the public key output by $\mathsf{KeyGen}(1^\lambda, 1^L)$. This follows from the semantic security of our Dual-Regev scheme which we proved in Theorem 14.

- **Certified deletion:** This ensures the following: Once the evaluation phase is completed and deletion phase is successful, the data $x \in \{0,1\}^\ell$ is *information-theoretically* deleted from the view of the server once a valid certificate is presented – provided the server is *honest* during the evaluation phase. This is essentially a consequence of Lemma 25 which guarantees that, once the evaluation phase is completed (and the server has performed the prescribed evaluation procedure), the leftover state is negligibly close to the original ciphertext. Therefore, the $\mathsf{EV\text{-}CD}$ security of our Dual-Regev scheme from Theorem 15 applies.

Therefore, our protocol only achieves a meaningful security guarantee in the so-called *semi-honest* setting [81], which requires that the adversary is honest during the evaluation phase of the protocol, but may later choose to maliciously analyze any leftover information which was collected throughout the protocol. A quantum analogue of the semi-honest adversarial model was also studied in [58] who consider so-called *specious* adversaries more generally. We leave the task of improving Protocol 1 as to allow for a possibly malicious server as an interesting open problem.

Chapter 5

REVOKING CRYPTOGRAPHIC KEYS

In this chapter, we build on the no-cloning principle of quantum mechanics and design cryptographic schemes with *key-revocation capabilities*. We consider several primitives with the guarantee that, once the secret key (represented as a quantum state) is successfully revoked from a user, they no longer have the ability to perform the same functionality as before. We define and construct several key-revocable cryptographic primitives; namely, pseudorandom functions, secret-key and public-key encryption, and even fully homomorphic encryption, assuming the quantum subexponential hardness of the learning with errors problem. Central to all of the constructions in this chapter is our key-revocable Dual-Regev public-key encryption scheme.

Organization. In Section 5.2, we show how to obtain a quantum discrete Gaussian sampler. This is a crucial subroutine which allows us to efficiently verify whether a given state corresponds to a particular Gaussian superposition. In Section 5.3 we prove the first quantum *Goldreich-Levin* theorem for large fields. Here, we rely on recent results on post-quantum reductions and *quantum rewinding*, which we also review in this section. Next, in Section 5.4, we describe the syntax of what a key-revocable public-key (as well as homomorphic) encryption scheme is. In Section 5.5, we introduce our key-revocable Dual-Regev scheme, which is the main section of this chapter. Our main technical result is a simultaneous search-to-decision reduction with quantum auxiliary input. In the subsequent section, we describe how to extend key-revocation capabilities towards homomorphic encryption schemes. Finally, in Section 5.7, we construct revocable *pseudorandom functions* by means of our key-revocable Dual-Regev scheme.

5.1 Introduction

Delegation and recovation of privilege are problems of great importance in cryptography. The problem of revocation in the context of digital signatures and certificates in the classical world is an especially thorny problem [121, 114]. As a motivating example, consider the setting of an employee at a company who takes a vacation and wishes to authorize a colleague to perform certain tasks on her behalf, tasks that involve handling sensitive data. Since the sensitive data is (required to be) encrypted, the employee must necessarily share her decryption keys with her colleague. When she returns from vacation, she would like to have her decryption key back; naturally, one would like to ensure that her colleague should not be able to decrypt future ciphertexts (which are encrypted under the same public key) once the key is "returned." Evidently, if the decryption key is a classical object, this is impossible to achieve.

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

Another, possibly non-obvious, application is to detecting malware attacks. Consider a malicious party who hacks into an electronic device and manages to steal a user's decryption keys. If cryptographic keys are represented by classical bits, it is inherently challenging to detect *phishing attacks* that compromise user keys. For all we know, the intruder could have stolen the user's decryption keys without leaving a trace. Indeed, a few years ago, decryption keys which were used to protect cell-phone communications [85] were successfully stolen by spies without being detected. With revocable cryptography, a malicious user successfully stealing a user key would invariably revoke the decryption capability from the user. This latter event can be detected.

Our Results in a Nutshell. We construct revocable cryptographic objects under standard cryptographic assumptions. This chapter features two main results. Our first main result constructs a key-revocable public-key encryption scheme, and our second main result constructs a key-revocable pseudorandom function. We obtain several corollaries and extensions, including key-revocable secret-key encryption and key-revocable fully homomorphic encryption. In all these primitives, secret keys are represented as quantum states that retain the functionality of the original secret keys. We design revocation procedures and guarantee that once a user successfully passes the procedure, they cannot compute the functionality any more. All our constructions are secure under the quantum subexponential hardness of learning with errors [112]—provided that revocation succeeds with high probability. At the heart of all of our contributions lies our result which shows that the Dual-Regev public-key encryption scheme of [68] satisfies revocation security.

Related Notions. There are several recent notions in quantum cryptography that are related to revocability. Of particular relevance is the stronger notion of copy-protection introduced by Aaronson [2]. Breaking the revocable security of a task gives the adversary a way to make two copies of a (possibly different) state both of which are capable of computing the same functionality. Thus, uncloneability is a stronger notion. However, the only known constructions of copy-protection [54, 96] rely on the heavy hammer of post-quantum secure indistinguishability obfuscation for which there are no known constructions based on well-studied assumptions. Our constructions, in contrast, rely on the post-quantum hardness of the standard learning with errors problem. Another related notion is the significantly weaker definition of secure software leasing [17] which guarantees that once the quantum state computing a functionality is returned, the *honest evaluation algorithm* cannot compute the original functionality. Yet another orthogonal notion is that of certifiably deleting *ciphertexts*, originating from the works of Unruh [127] and Broadbent and Islam [40]. In contrast, our goal is to delegate and revoke *cryptographic capabilities* enabled by private keys. For detailed comparisons, we refer the reader to Section 5.1.

Our Contributions in More Detail

We present our results in more detail below. First, we introduce the notion of key-revocable public-key encryption. Our main result is that dual-Regev public-key encryption scheme [68] satisfies revocation security. After that, we study revocation security in the context of fully homomorphic encryption and pseudorandom functions.

Key-Revocable Public-Key Encryption. We consider public-key encryption schemes where the decryption key, modeled as a quantum state, can be delegated to a third party and can later be revoked [70]. The syntax of a key-revocable public-key scheme (Definition 39) is as follows:

- $\mathsf{KeyGen}(1^\lambda)$: this is a setup procedure which outputs a public key pk, a master secret key msk and a decryption key ϱ_{sk}. While the master secret key is typically a classical string, the decryption key is modeled as a quantum state. (The use cases of msk and ϱ_{sk} are different, as will be evident below.)
- $\mathsf{Enc}(\mathsf{pk}, x)$: this is the regular classical encryption algorithm which outputs a ciphertext CT.
- $\mathsf{Dec}(\varrho_{\mathsf{sk}}, \mathsf{CT})$: this is a quantum algorithm which takes as input the quantum decryption key sk and a classical ciphertext, and produces a plaintext.
- $\mathsf{Revoke}(\mathsf{pk}, \mathsf{msk}, \sigma)$: this is the revocation procedure that outputs Valid or $\mathsf{Invalid}$. If σ equals the decryption key sk, then Revoke is expected to output Valid with high probability.

After the decryption key is returned, we require that the sender loses its ability to decrypt ciphertexts. This is formalized as follows (see Definition 40): conditioned on revocation being successful, the adversary should not be able to distinguish whether it is given an encryption of a message versus uniform distribution over the ciphertext space with more than negligible advantage. Moreover, we require that revocation succeeds with a probability negligibly close to 1 (more on this later).

We prove the following in Theorem 24.

Theorem (Informal)**.** *Assuming that the* LWE *and* SIS *problems with subexponential modulus are hard against quantum adversaries running in subexponential time (see Section 2.6), there exists a key-revocable public-key encryption scheme.*

Due to the quantum reduction from SIS to LWE [120], the two assumptions are, in some sense, equivalent. Therefore, we can in principle rely on the subexponential hardness of LWE alone.

Our results improve upon prior works, which either use post-quantum secure indistinguishability obfuscation [70, 54] or consider the weaker private-key setting [88].

Key-Revocable Fully Homomorphic Encryption. We go beyond the traditional public-key setting and design the first *fully homomorphic encryption* (FHE) scheme [67, 37] with key-revocation capabilities. Our construction is based on a variant of the (leveled) FHE scheme of Gentry, Sahai and Waters [69], which we extend to a key-revocable encryption scheme using Gaussian superpositions. The syntax of a key-revocable FHE scheme is the same as in the key-revocable public-key setting from before (Definition 39), except for the additional algorithm Eval which is the same as in a regular FHE scheme. We prove the following in Theorem 33.

Theorem (Informal)**.** *Assuming that the* LWE *and* SIS *problems with subexponential modulus are hard against quantum adversaries running in subexponential time (see Section 2.6), there exists a key-revocable (leveled) fully homomorphic encryption scheme.*

We prove the theorem by invoking the security of our key-revocable Dual-Regev public-key encryption scheme in Section 5.5.

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

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

After the quantum key ϱ_k is successfully returned, we require that the sender loses its ability to evaluate the PRF. This is formalized as follows (see Definition 44): any efficient adversary can simultaneously pass the revocation phase and succeed in predicting the output of a pseudorandom function on a challenge input x^* versus uniform with advantage at most $\mathsf{negl}(\lambda)$. In fact, we consider a more general definition where the adversary receives many challenge inputs instead of just one challenge input.

We give the first construction of key-revocable pseudorandom functions (PRFs) from standard assumptions. Previous schemes implicit in [54] either require indistinguishability obfuscation, or considered weaker notions of revocable PRFs in the form of *secure software leasing* [17, 89], which merely prevents the possiblity of *honestly* evaluating the PRF once the key is revoked.

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

Theorem (Informal)**.** *Assuming that the* LWE *and* SIS *problems with subexponential modulus are hard against quantum adversaries running in subexponential time (see Section 2.6), there exist key-revocable pseudorandom functions.*

Revocable pseudorandom functions immediately give us key-revocable (many time secure) secret-key encrypton schemes.

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

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

We show how to achieve all of our results based on a conjecture, that we call *Simultaneous Dual-Regev Extraction* (SDRE) conjecture. Informally, the conjecture states that if Dual-Regev PKE is not key-revocable then there exists a QPT adversary who given a Gaussian superposition $|\psi_{\mathbf{y}}\rangle$ of short preimages mapping a random matrix $\mathbf{A}$ to a vector $\mathbf{y}$ can simultaneously produce $|\psi_{\mathbf{y}}\rangle$ and a short vector in the support of $|\psi_{\mathbf{y}}\rangle$ with non-negligible probability.

In more detail, the SDRE conjecture states the following: suppose the Dual-Regev PKE is not key-revocable (according to the above stronger definition) then there exists a QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that the following holds:

- $\mathcal{A}_1$ is given $(\mathbf{A}, \mathbf{y}, |\psi_\mathbf{y}\rangle)$, where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ and $|\psi_\mathbf{y}\rangle$ is a Gaussian superposition of all the short vectors mapping $\mathbf{A}$ to $\mathbf{y}$. It produces a bipartite state on two registers R and AUX.
- A projective measurement on R is applied that projects onto the state $|\psi_\mathbf{y}\rangle\langle\psi_\mathbf{y}|$, and $\mathcal{A}_2$ is run on AUX. We require that the (simultaneous) probability that the projective measurement succeeds and $\mathcal{A}_2$ outputs a short preimage mapping $\mathbf{A}$ to $\mathbf{y}$ should be inverse polynomial.

The difficulty in proving the conjecture lies in the fact that one needs to invoke the LWE assumption with respect to $\mathcal{A}_2$ who holds AUX, while at the same time guaranteeing that an inefficient projective measurement succeeds on a separate register R. We leave proving (or refuting) the above conjecture to future works.

Discussion: Unclonable Cryptography from LWE. Over the years, the existence of many fundamental cryptographic primitives such as pseudorandom functions [21], fully homomorphic encryption [37], attribute-based encryption [33] and succinct argument systems [50] have been based on the existence of learning with errors. In fact, as far as we know, there are only a few foundational primitives remaining (indistinguishability obfuscation is one such example) whose existence is not (yet) known to be based on learning with errors.

This situation is quite different in the world of unclonable cryptography. Most of the prominent results have information-theoretic guarantees but restricted functionalities [40, 41] or are based on the existence of post-quantum indistinguishability obfuscation [137, 54]. While there are works [90] that do propose lattice-based constructions of unclonable primitives, there are still many primitives, such as quantum money and quantum copy-protection, whose feasibility we would like to establish based on the existence of learning with errors. We hope that our work presents new toolkits towards building more unclonable primitives from LWE.

Independent and Concurrent Work. Independently and concurrently, Agrawal et al. [8] explored the notion of public-key encryption with secure leasing which is related to key-revocable public-key encryption. They achieve a generic construction based on any post-quantum secure public-key encryption whereas our notion is based on the post-quantum hardness of learning with errors. They also explore other notions of advanced encryption with secure leasing including attribute-based encryption and functional encryption, which are not explored in our work.

On the other hand, their construction of revocable public-key encryption involves many abstractions whereas our construction is based on the versatile Dual-Regev public-key encryption scheme. Additionally, we obtain key-revocable *fully homomorphic encryption* and key-revocable *pseudorandom functions* which are unique to our work.

- *Advanced notions*: We obtain key-revocable *fully homomorphic encryption* and key-revocable *pseudorandom functions* which are unique to our work. They explore other notions of advanced encryption with secure leasing including attribute-based encryption and functional encryption, which are not explored in our work.
- *Public-key encryption*: They achieve a generic construction based on any post-quantum secure public-key encryption[1] whereas our notion is based on the post-quantum hardness of the learning with errors problem or the SDRE conjecture. Their construction of revocable public-key encryption involves many complex abstractions whereas our construction is based on the versatile Dual-Regev public-key encryption scheme.

Overview

We now give a technical overview of our constructions and their high level proof ideas. We begin with the key-revocable public-key encryption construction. A natural idea would be to start with Regev's public-key encryption scheme [112] and to then upgrade the construction in order to make it revocable. However, natural attempts to associate an unclonable quantum state with the decryption key fail and thus, we instead consider the Dual-Regev public-key encryption scheme and make it key-revocable. We describe the scheme below.

Key-Revocable Dual-Regev Public-Key Encryption. Our first construction is based on the *Dual-Regev* public-key encryption scheme [68] and makes use of Gaussian superpositions which serve as a quantum decryption key. We give an overview of Construction 8 below.

- $\mathsf{KeyGen}(1^n)$: sample a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ along with a *short trapdoor basis* $\mathsf{td}_{\mathbf{A}}$. To generate the decryption key, we employ the following procedure[2]: Using the matrix $\mathbf{A}$ as input, first create a Gaussian superposition of short vectors in $\mathbb{Z}^m \cap (-\frac{q}{2}, \frac{q}{2}]^m$, denoted by[3]

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

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

[2] In Section 5.2, this is formalized as the procedure $\mathsf{GenGauss}$ (see Algorithm 3).

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

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

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

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

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

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

- $\mathsf{Dec}(|\mathsf{sk}\rangle, \mathsf{CT})$: to decrypt a ciphertext CT using the decryption key $\mathsf{sk} = |\psi_{\mathbf{y}}\rangle$, first apply the unitary $U : |\mathbf{x}\rangle|0\rangle \to |\mathbf{x}\rangle|\mathsf{CT}\cdot(-\mathbf{x}, 1)^\intercal\rangle$ on input $|\psi_{\mathbf{y}}\rangle|0\rangle$, and then measure the second register in the computational basis. Because $|\psi_{\mathbf{y}}\rangle$ is a superposition of short vectors $\mathbf{x}$ subject to $\mathbf{A}\cdot\mathbf{x} = \mathbf{y} \pmod q$, we obtain an approximation of $\mu \cdot \lfloor \frac{q}{2} \rfloor$ from which we can recover μ.[4]

- $\mathsf{Revoke}(\mathsf{pk}, \mathsf{msk}, \varrho)$: to verify the returned state ϱ given as input the public key $(\mathbf{A}, \mathbf{y})$ and master secret key $\mathsf{td}_{\mathbf{A}}$, apply the projective measurement $\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, \mathbb{1} - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\}$ onto ϱ. Output Valid, if the measurement succeeds, and output $\mathsf{Invalid}$, otherwise.

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

One of our contributions is to design a *quantum discrete Gaussian sampler for q-ary lattices*[5] which, given as input $(\mathbf{A}, \mathbf{y}, \mathsf{td}_{\mathbf{A}}, \sigma)$, implements a unitary that efficiently prepares the Gaussian superposition $|\psi_{\mathbf{y}}\rangle$ from scratch with access to the trapdoor $\mathsf{td}_{\mathbf{A}}$. At a high level, our Gaussian sampler can be alternately thought of as an explicit quantum reduction from the *inhomogenous* SIS problem [9] to the search variant of the LWE problem (see Section 5.2).

[4] For approriate choices of parameters, decryption via rounding succeeds at outputting μ with overwhelming probability and hence we can invoke the "*Almost as Good as New Lemma*" [3] to recover the original state $|\psi_{\mathbf{y}}\rangle$.

[5] In Section 5.2, this is formalized as the procedure QSampGauss (see Algorithm 4).

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

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

- Firstly, measuring the returned state should give a vector different from $\mathbf{x}$ with non-negligible probability. In order to prove this, we need to argue that the squared ampltidue of every term is bounded away from 1. We prove this statement (Lemma 17) holds as long as $\mathbf{A}$ is full rank.
- Secondly, the reduction to SIS would only get as input $\mathbf{A}$ and not a trapdoor for $\mathbf{A}$. This means that it will no longer be possible for the reduction to actually check whether the state returned by the adversary is valid. We observe that, instead of first verifying whether the returned state is valid and then measuring in the computational basis, we can in fact skip verification and immediately go ahead and measure the state in the computational basis; this is implicit in the analysis in the proof of Lemma 29.
- Finally, the adversary could have entangled the returned state with its residual state in such a way that measuring the returned state always yields the same vector $\mathbf{x}$ as the one extracted from the adversary. In the same analysis in the proof of Lemma 29, we prove that, even if the adversary entangles its state with the returned state, with non-negligible probability we get two distinct short vectors mapping $\mathbf{A}$ to $\mathbf{y}$.

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

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

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

Ignoring register R: For now, let us ignore the fact that Revoke is simultaneously applied on system R. Inspired by techniques from the *leakage resilience* literature [57], we now make the following observation. Letting $\mathbf{y} = \mathbf{A} \cdot \mathbf{x}_0 \pmod{q}$, for some Gaussian vector $\mathbf{x}_0$ with distribution proportional to $\varrho_\sigma(\mathbf{x}_0)$, the former sample can be written as $(\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal, (\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal) \cdot \mathbf{x}_0 + e')$. Here, we assume a *noise flooding* regime in which the noise magnitude of e' is significantly larger than that of $\mathbf{e}^\intercal \cdot \mathbf{x}_0$. Because the distributions are statistically close, the distinguisher $\mathcal{D}$ must succeed at distinguishing the sample from uniform with probability negligibly close to ϵ. Finally, we invoke the LWE assumption and claim that the same distinguishing advantage persists, even if we replace $(\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal)$ with a random string $\mathbf{u} \in \mathbb{Z}_q^m$. Here, we rely on the fact that the underlying LWE sample is, in some sense, independent of the auxiliary input Aux handed to the distinguisher $\mathcal{D}$. To show that this is the case, we need to argue that the reduction can generate the appropriate inputs to $\mathcal{D}$ on input $\mathbf{A}$; in particular it should be able to generate the auxiliary input Aux (which depends on a state $|\psi_\mathbf{y}\rangle$), while simultaneously producing a Gaussian vector $\mathbf{x}_0$ such that $\mathbf{A} \cdot \mathbf{x}_0 = \mathbf{y} \pmod{q}$. Note that this seems to violate the SIS assumption, since the ability to produce both a superposition $|\psi_\mathbf{y}\rangle$ of pre-images and a single pre-image $\mathbf{x}_0$ would allow one to obtain a collision for $\mathbf{y}$.

Invoking Gaussian-collapsing: To overcome this issue, we ask the reduction to generate the quantum auxiliary input in a different way; rather than computing Aux as a function of $|\psi_\mathbf{y}\rangle$, we compute it as a function of $|\mathbf{x}_0\rangle$, where $\mathbf{x}_0$ results from *collapsing* the state $|\psi_\mathbf{y}\rangle$ via a measurement in the computational basis. By invoking the *Gaussian collapsing property* [109], we can show that the auxiliary information computed using $|\psi_\mathbf{y}\rangle$ is computationally indistinguishable from the auxiliary information computed using $|\mathbf{x}_0\rangle$. Once we invoke the collapsed version of $|\psi_\mathbf{y}\rangle$, we can carry out the reduction and conclude that $\mathcal{D}$ can distinguish between the samples $(\mathbf{u}, \mathbf{u}^\intercal \mathbf{x}_0)$ and $(\mathbf{u}, r)$, where $\mathbf{u}$ and r are random and $\mathbf{x}_0$ is Gaussian, with advantage negligibly close to ϵ.[6] Notice that $\mathcal{D}$ now resembles a so-called *Goldreich-Levin* distinguisher [71].

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

Reduction to Goldreich-Levin: Assuming the existence of a quantum Goldreich-Levin theorem for the field $\mathbb{Z}_q$, one could then convert $\mathcal{D}$ into an extractor that extracts $\mathbf{x}_0$ with high probability. Prior to our work, a quantum Goldreich-Levin theorem was only known for $\mathbb{Z}_2$ [7, 54]. In particular, it is unclear how to extend prior work towards higher order fields $\mathbb{Z}_q$ because the interference pattern in the analysis of the quantum extractor does not seem to generalize beyond the case when $q = 2$. Fortunately, we can rely on the *classical* Goldreich-Levin theorem for finite fields due to Dodis et al. [57], as well as recent work by Bitansky, Brakerski and Kalai. [32] which shows that a large class of classical reductions can be generically converted into a quantum reductions. This allows us to obtain the first quantum Goldreich-Levin theorem for large fields, which we prove in Section 5.3. Specifically, we can show that a distinguisher $\mathcal{D}$ that, given auxiliary input $\textsc{Aux}$, can distinguish between $(\mathbf{u}, \mathbf{u}^\intercal \mathbf{x}_0)$ and $(\mathbf{u}, r)$ with advantage ε can be converted into a quantum extractor that can extract $\mathbf{x}_0$ given $\textsc{Aux}$ in time $\mathsf{poly}(1/\varepsilon, q)$ with probability $\mathsf{poly}(\varepsilon, 1/q)$.

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

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

Applications

We leverage our result of key-revocable Dual-Regev encryption to get key-revocable fully homomorphic encryption and revocable pseudorandom functions.

Key-Revocable Dual-Regev Fully Homomorphic Encryption. Our first application of our key-revocable public-key encryption concerns fully homomorphic encryption schemes. We extend our key-revocable Dual-Regev scheme towards a (leveled) FHE scheme in Construction 9 by using the $\mathsf{DualGSW}$ variant of the FHE scheme by Gentry, Sahai, and Waters [69, 99].

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

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

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

Note that the DualGSW ciphertext can be thought of as a column-wise concatenation of N-many independent Dual-Regev ciphertexts. In Theorem 33, we prove the security of our construction by invoking the security of our key-revocable Dual-Regev scheme.

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

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

1. First show that there exists a μ-revocable pseudorandom function for $\mu = 1$. Here, μ-revocation security means the adversary receives μ-many random inputs after revocation.
2. Next, we show that any 1-revocable pseudorandom function also satisfies the stronger notion of revocation security where there is no a priori bound on the number of challenge inputs received by the adversary.
3. Finally, we show that we can generically upgrade any revocable PRF in such a way that it also satisfies the traditional pseudorandomness property.

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

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

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

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

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

- For every $i \in [\ell]$, include $\mathbf{S}_{i,b}\mathbf{A} + \mathbf{E}_{i,b}$ in ϱ_k, where $i \in [\ell]$ and $b \in \{0, 1\}$. We sample $\mathbf{S}_{i,b}$ and $\mathbf{E}_{i,b}$ from a discrete Gaussian distribution with appropriate standard deviation $\sigma > 0$.
- Include $|\psi_{\mathbf{y}}\rangle$ which, as defined in the key-revocable Dual-Regev construction, is a Gaussian superposition of short solutions mapping $\mathbf{A}$ to $\mathbf{y}$.

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

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

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

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

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

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

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

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

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

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

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

Next, we modify ϱ_k as follows: instead of generating, $\mathbf{S}_b^{(i,j,\tau)}\mathbf{A} + \mathbf{E}_b^{(i,j,\tau)}$, we instead generate $\mathbf{S}_b^{(i,j,\tau)}\mathbf{A} + \mathbf{E}_b^{(i,j,\tau)} + \mathsf{M}_b^{(i,j,k)}$, for any set of matrices $\{\mathsf{M}_b^{(i,j,\tau)}\}$. The change should be undetectable to a computationally bounded adversary, thanks to the quantum hardness of learning with errors. In the security proof, we set up the challenge input x^* in such a way that summing up the matrices $\mathsf{M}_{x_i^*}^{(i,j,\tau)}$ corresponds to CT_1^*, where CT_1^* is part of the key-revocable Dual-Regev challenge ciphertext as discussed above. With this modification, when ϱ_k is evaluated on x^*, we get an output that is close to $\mathsf{CT}_2^* + \mathbf{u}$, where $\mathbf{u} \approx \sum_{i\in[n],j\in[m],\tau\in[\lceil\log(q)\rceil]} \mathbf{y}$ is known to the reduction (discussed above)—thereby violating the security of key-revocable Dual-Regev scheme.

Related Work

Let us now briefly mention related work.

Copy-Protection. Of particular relevance to our work is the foundational notion of copy-protection introduced by Aaronson [2]. Informally speaking, a copy-protection scheme is a compiler that transforms programs into quantum states in such a way that using the resulting states, one can run the original program. Yet, the security guarantee stipulates that any adversary given one copy of the state cannot produce a bipartite state wherein both parts compute the original program.

While copy-protection is known to be impossible for arbitrary unlearnable functions [17, 13], identifying interesting functionalities for which copy-protection is feasible has been an active research direction [52, 18, 16]. Of particular significance is the problem of copy-protecting cryptographic functionalities, such as decryption and signing functionalities. Coladangelo et al. [54] took the first step in this direction and showed that it is feasible to copy-protect decryption functionalities and pseudorandom functions assuming the existence of post-quantum indistinguishability obfuscation. While a very significant first step, the assumption of post-quantum iO is unsatisfactory: there have been numerous post-quantum iO candidate proposals [19, 49, 39, 55, 65, 130], but not one of them have been based on well-studied assumptions[7].

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

[7] We remark that, there do exist post-quantum-insecure iO schemes based on well-founded assumptions [86].

Secure Software Leasing. Another primitive relavent to revocable cryptography is secure software leasing [17]. The notion of secure software leasing states that any program can be compiled into a functionally equivalent program, represented as a quantum state, in such a way that once the compiled program is returned back[8], the honest evaluation algorithm on the residual state cannot compute the original functionality. Key-revocable encryption can be viewed as secure software leasing for decryption algorithms. However, unlike secure software leasing, key-revocable encryption satisfies a much stronger security guarantee, where there is no restriction on the adversary to run honestly after returning back the software. Secure leasing for different functionalities, namely, point functions [52, 43], evasive functions [17, 90] and pseudorandom functions [6] have been studied by recent works.

Encryption Schemes with Revocable Ciphertexts. Unruh [127] proposed a (private-key) quantum timed-release encryption scheme that is *revocable*, i.e., it allows a user to *return* the ciphertext of a quantum timed-release encryption scheme, thereby losing all access to the data. Unruh's scheme uses conjugate coding [132, 29] and relies on the *monogamy of entanglement* in order to guarantee that revocation necessarily erases information about the plaintext. Broadbent and Islam [40] introduced the notion of *certified deletion*[9] and constructed a private-key quantum encryption scheme with the aforementioned feature which is inspired by the quantum key distribution protocol [29, 122]. In contrast with Unruh's [127] notion of revocable quantum ciphertexts which are eventually returned and verified, Broadbent and Islam [40] consider certificates which are entirely classical. Moreover, the security definition requires that, once the certificate is successfully verified, the plaintext remains hidden even if the secret key is later revealed.

Using a hybrid encryption scheme, Hiroka, Morimae, Nishimaki, and Yamakawa [83] extended the scheme in [41] to both public-key and attribute-based encryption with certified deletion via *receiver non-committing* encryption [87, 45]. 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. Bartusek and Khurana [23] revisited the notion of certified deletion and presented a unified approach for how to generically convert any public-key, attribute-based, fully-homomorphic, timed-release or witness encryption scheme into an equivalent quantum encryption scheme with certified deletion. In particular, they considered a stronger notion called *certified everlasting security* which allows the adversary to be computationally unbounded once a valid deletion certificate is submitted.

[8] According to the terminology of [17], this refers to finite term secure software leasing.

[9] This notion is incomparable with another related notion called unclonable encryption [41, 15, 18], which informally guarantees that it should be infeasible to clone quantum ciphertexts without losing information about the encrypted message.

5.2 Quantum Discrete Gaussian Sampling for q-ary Lattices

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

Gaussian Superpositions

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

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

Note that the state $|\psi\rangle$ is not normalized for convenience and ease of notation. The tail bound in Lemma 11 implies that (the normalized variant of) $|\psi\rangle$ is within negligible trace distance of a *truncated* discrete Gaussian superposition $|\tilde{\psi}\rangle$ with support $\{\mathbf{x} \in \mathbb{Z}_q^m : \|\mathbf{x}\| \leq \sigma\sqrt{\frac{m}{2}}\}$, where

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

In this work, we consider Gaussian superpositions with parameter $\sigma = \Omega(\sqrt{m})$ which can be efficiently implemented using standard quantum state preparation techniques; for example using *quantum rejection sampling* and the *Grover-Rudolph algorithm* [78, 112, 36, 38].

Gaussian coset states. Our key-revocable encryption schemes in Section 5.5 and Section 5.6 rely on Gaussian superpositions over $\mathbf{x} \in \mathbb{Z}_q^m$ subject to a constraint of the form $\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod q$, for some matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and image $\mathbf{y} \in \mathbb{Z}_q^n$. In Algorithm 3, we give a procedure called $\mathsf{GenGauss}$ that, on input $\mathbf{A}$ and $\sigma > 0$, generates a Gaussian superposition state of the form

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

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

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

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

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

Gaussian-collapsing hash functions. We recall the following variant of the Gaussian-collapsing property of the Ajtai hash function which we previously showed in Theorem 9.

Corollary 3 (Gaussian-collapsing property)**.** *Let $n \in \mathbb{N}$ and q be a prime with $m \geq 2n\log q$, each parameterized by $\lambda \in \mathbb{N}$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$. Then, the following samples are computationally indistinguishable assuming the quantum hardness of decisional $\mathsf{LWE}^m_{n,q,\alpha q}$, for any noise ratio $\alpha \in (0,1)$ with relative noise magnitude $1/\alpha = \sigma \cdot 2^{o(n)}$:*

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

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

Algorithm: GenGauss

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

```
Algorithm 3: GenGauss($\mathbf{A},\sigma$)

Input: Matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and parameter $\sigma = \Omega(\sqrt{m})$.
Output: Gaussian state $|\psi_{\mathbf{y}}\rangle$ and $\mathbf{y} \in \mathbb{Z}_q^n$.

Prepare a Gaussian superposition in system $X$ with parameter $\sigma > 0$:
    $|\psi\rangle_{XY} = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \varrho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle_X \otimes |\mathbf{0}\rangle_Y .$
Apply the unitary $U_{\mathbf{A}} : |\mathbf{x}\rangle\,|\mathbf{0}\rangle \rightarrow |\mathbf{x}\rangle\,|\mathbf{A}\cdot\mathbf{x} \pmod{q}\rangle$ on systems $X$ and $Y$:
    $|\psi\rangle_{XY} = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \varrho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle_X \otimes |\mathbf{A}\cdot\mathbf{x} \pmod{q}\rangle_Y .$
Measure system $Y$ in the computational basis, resulting in the state
    $|\psi_{\mathbf{y}}\rangle_{XY} = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{Ax}=\mathbf{y}}} \varrho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle_X \otimes |\mathbf{y}\rangle_Y .$
Output the state $|\psi_{\mathbf{y}}\rangle$ in system $X$ and the outcome $\mathbf{y} \in \mathbb{Z}_q^n$ in system $Y$.
```

Algorithm: QSampGauss

Recall that, in Algorithm 3, we gave a procedure called $\mathsf{GenGauss}(\mathbf{A},\sigma)$ that prepares a Gaussian coset state $|\psi_{\mathbf{y}}\rangle$, for a randomly generated $\mathbf{y} \in \mathbb{Z}_q^n$. In general, however, generating a specific Gaussian coset state on input $(\mathbf{A},\mathbf{y})$ requires a *short trapdoor basis* $\mathsf{td}_{\mathbf{A}}$ for the matrix $\mathbf{A}$. This task can be thought of as a quantum analogue of the *discrete Gaussian sampling problem* [68], where the goal is to output a sample $\mathbf{x} \sim D_{\mathbb{Z}^m,\sigma}$ such that $\mathbf{A}\cdot\mathbf{x} = \mathbf{y} \pmod q$ on input $(\mathbf{A},\mathbf{y})$ and $\sigma > 0$.

In Algorithm 4, we give a procedure called $\mathsf{QSampGauss}$ which, on input $(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ generates a specific Gaussian coset state $|\psi_{\mathbf{y}}\rangle$ of the form

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

Our procedure $\mathsf{QSampGauss}$ in Algorithm 4 can be thought of as an explicit quantum reduction from $\mathsf{ISIS}^m_{n,q,\sigma\sqrt{m/2}}$ to $\mathsf{LWE}^m_{n,q,q/\sqrt{2}\sigma}$ which is inspired by the quantum reduction of Stehlé et al. [120] which reduces SIS to LWE. To obtain the aforementioned reduction, one simply needs to replace the procedure $\mathsf{Invert}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \cdot)$ in Step 4 in Algorithm 4 with a solver for the $\mathsf{LWE}^m_{n,q,q/\sqrt{2}\sigma}$ problem.

In Theorem 16, we prove the correctness of Algorithm 4. As a technical ingredient, we rely on a *duality lemma* from Lemma 20 that characterizes the Fourier transform of a Gaussian coset state in terms of its dual state. Note that $|\psi_{\mathbf{y}}\rangle$ corresponds to a Gaussian superposition over a lattice coset,

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

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

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

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

```
Algorithm 4: QSampGauss($\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma$)
Input: Matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a trapdoor $\mathsf{td}_{\mathbf{A}}$, an image $\mathbf{y} \in \mathbb{Z}_q^n$ and parameter $\sigma = O(\frac{q}{\sqrt{m}})$.
Output: Gaussian state $|\psi_{\mathbf{y}}\rangle$.
Prepare the following superposition with parameter $q/\sigma > 0$:
    $$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} |\mathbf{s}\rangle \otimes \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \varrho_{q/\sigma}(\mathbf{e}) |\mathbf{e}\rangle \otimes |\mathbf{0}\rangle$$
Apply the generalized Pauli operator $\mathbf{Z}_q^{-\mathbf{y}}$ on the first register, resulting in the state
    $$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle} |\mathbf{s}\rangle \otimes \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \varrho_{q/\sigma}(\mathbf{e}) |\mathbf{e}\rangle \otimes |\mathbf{0}\rangle$$
Apply the unitary $U_{\mathbf{A}} : |\mathbf{s}\rangle |\mathbf{e}\rangle |\mathbf{0}\rangle \to |\mathbf{s}\rangle |\mathbf{e}\rangle |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod{q}\rangle$, resulting in the state
    $$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \varrho_{q/\sigma}(\mathbf{e}) \, \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle} |\mathbf{s}\rangle |\mathbf{e}\rangle |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod{q}\rangle$$
Coherently run Invert($\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \cdot$) on the third register in order to uncompute the first and the
  second register, resulting in a state that is close in trace distance to the following state:
    $$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{\mathbf{e} \in \mathbb{Z}_q^m} \varrho_{q/\sigma}(\mathbf{e}) \, \omega_q^{-\langle \mathbf{s}, \mathbf{y} \rangle} |0\rangle |0\rangle |\mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal \pmod{q}\rangle$$
Discard the first two registers. Apply the (inverse) quantum Fourier transform and output the
  resulting state.
```

Let us now prove the correctness of Algorithm 4.

Theorem 16 (Quantum Discrete Gaussian Sampler)**.** *Let $n \in \mathbb{N}$, q be a prime with $m \geq 2n \log q$ and $\sqrt{8m} < \sigma < q/\sqrt{8m}$. Let $(\mathbf{A}, \mathsf{td}_{\mathbf{A}}) \leftarrow$* GenTrap$(1^n, 1^m, q)$ *be sampled as in Theorem 2 and let $\mathbf{y} \in \mathbb{Z}_q^n$ be arbitrary. Then, with overwhelming probability,* QSampGauss$(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ *in*

Algorithm 4 outputs a state which is within negligible trace distance of the (normalized variant of the) state,

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

Proof. From Lemma 8 and Theorem 2, it follows that $(\mathbf{A},\mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n,1^m,q)$ yields a matrix $\mathbf{A}\in\mathbb{Z}_q^{n\times m}$ be a matrix whose columns generate $\mathbb{Z}_q^n$ with overwhelming probability. Moreover, since $\sqrt{8m} < \sigma < q/\sqrt{8m}$, the inversion procedure $\mathsf{Invert}(\mathbf{A},\mathsf{td}_{\mathbf{A}},\cdot)$ from Theorem 2 in Step 4 in Algorithm 4 succeeds with overwhelming probability at generating the Gaussian state

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

Applying the (inverse) quantum Fourier transform $\mathsf{FT}_q^\dagger$, the claim then follows from Lemma 20. □

5.3 Quantum Goldreich-Levin Theorem for Large Fields

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

Post-Quantum Reductions and Quantum Rewinding

We first review some recent work by Bitansky, Brakerski and Kalai [32] that enables us to convert a wide range of classical reductions into post-quantum reductions (which allow for quantum auxiliary input) in a constructive manner. We first review some basic terminology from [32].

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

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

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

We use the following theorem.

Theorem 17 ([32], adapted from Theorem 7.1)**.** *Let $c \in \mathbb{R}$. Suppose that there exists a classical reduction from solving a non-interactive assumption Q to solving a non-interactive assumption P such that the following holds: if the P-solver has advantage $\varepsilon > 0$ then the Q-solver has advantage c (independent of ε) with running time $\mathsf{poly}(1/\varepsilon, c, \lambda)$.*

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

Remark 18. *We note that [32] consider a more general theorem where the advantage of the classical Q-solver can depend on the advantage of the P-solver. But in the case when the classical Q-solver's advantage is independent of the P-solver's advantage then, as reflected in the above theorem, it turns out the advantage of the quantum Q-solver is the same as the classical Q-solver.*

Goldreich-Levin Theorems for Large Fields

The following result is implicit in the work of Dodis et al. [57].

Theorem 19 (Classical Goldreich-Levin Theorem for Finite Fields, [57], Theorem 1)**.** *Let q be a prime and $m \in \mathbb{N}$. Let $H = \{\mathbf{x} \in \mathbb{Z}_q^m \ : \ \|\mathbf{x}\| \leq \sigma\sqrt{m}\}$ be a subset of $\mathbb{Z}_q^m$, for some $\sigma > 0$. Let $f : H \to \{0,1\}^*$ be any (possibly randomized) function. Suppose there exists a distinguisher $\mathcal{D}$ that runs in time $T(\mathcal{D})$ and has the property that*

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

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

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

Using the constructive post-quantum reduction from Theorem 17, we can convert Theorem 19 into a quantum Goldreich-Levin Theorem for finite fields, and obtain the following.

Theorem 20 (Quantum Goldreich-Levin Theorem for Large Fields)**.** *Let q be a prime and $m \in \mathbb{N}$. Let $H = \{\mathbf{x} \in \mathbb{Z}_q^m \ : \ \|\mathbf{x}\| \leq \sigma\sqrt{m}\}$ be a subset of $\mathbb{Z}_q^m$, for some $\sigma > 0$. Let $\Phi : \mathcal{L}(\mathcal{H}_q^m) \to \mathcal{L}(\mathcal{H}_{AUX})$ be any CPTP map with auxiliary system $\mathcal{H}_{AUX}$. Suppose there exists a quantum distinguisher $\mathcal{D}$ that runs in time $T(\mathcal{D})$ and has the property that*

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

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

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

Proof. The proof follows immediately by combining Theorem 19 and Theorem 17. □

5.4 Definition: Key-Revocable Public-Key Encryption

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

Definition 39 (Key-Revocable Public-Key Encryption)**.** *A key-revocable public-key encryption scheme consists efficient algorithms* $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$*, where* Enc *is a* PPT *algorithm and* $\mathsf{KeyGen}, \mathsf{Dec}$ *and* Revoke *are* QPT *algorithms defined as follows:*

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

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

$$\Pr\left[x \leftarrow \mathsf{Dec}(\varrho_{\mathsf{sk}}, \mathsf{CT}) \ : \ \substack{(\mathsf{pk}, \mathsf{msk}, \varrho_{\mathsf{sk}}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{pk}, x)}\right] \geq 1 - \nu(\lambda),$$

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

Correctness of Revocation. The following holds:

$$\Pr\left[\mathsf{Valid} \leftarrow \mathsf{Revoke}(\mathsf{pk}, \mathsf{msk}, \varrho_{\mathsf{sk}}) \ : \ (\mathsf{pk}, \mathsf{msk}, \varrho_{\mathsf{sk}}) \leftarrow \mathsf{KeyGen}(1^\lambda)\right] \geq 1 - \nu(\lambda),$$

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

Remark 21. *Using the "Almost As Good As New" Lemma (Lemma 1), the procedure* Dec *can easily be purified to obtain another quantum circuit* $\widetilde{\mathsf{Dec}}$ *such that* $\widetilde{\mathsf{Dec}}(\varrho_{\mathsf{sk}}, \mathsf{CT})$ *yields $(x, \varrho_{\mathsf{sk}}')$ with probability at least $1 - \nu(\lambda)$ and, moreover,* CT *is an encryption of x and $\|\varrho_{\mathsf{sk}}' - \varrho_{\mathsf{sk}}\|_{\mathrm{tr}} \leq \nu'(\lambda)$, where $\nu'(\lambda)$ is another negligible function.*

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

Initialization Phase:

- The challenger runs $(\mathsf{pk}, \mathsf{msk}, \varrho_{\mathsf{sk}}) \leftarrow \mathsf{KeyGen}(1^{\lambda})$ and sends $(\mathsf{pk}, \varrho_{\mathsf{sk}})$ to $\mathcal{A}$.

Revocation Phase:

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

Guessing Phase:

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

Figure 5.1: Security Experiment.

Security Definition

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

Definition 40. *A key-revocable public-key encryption scheme* $\Sigma = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ *is* (ϵ, δ)*-secure if, for every* QPT *adversary* $\mathcal{A}$ *with*

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

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

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

where $\mathsf{Expt}_{\Sigma,\mathcal{A}}(1^{\lambda}, b)$ *is as defined in Figure 5.1. If* $\delta(\lambda) = 1 - 1/\mathsf{poly}(\lambda)$ *and* $\varepsilon(\lambda) = \mathsf{negl}(\lambda)$*, we simply say the key-revocable public-key encryption scheme is secure.*

Remark 22. *Our security definition is similar to the one proposed by Agrawal et al. [8] in the context of public-key encryption with secure leasing.*

Key-Revocable Public-Key Fully Homomorphic Encryption

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

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

Remark 23. *Sometimes we allow* KeyGen *to additionally take as input different parameters associated with the implementations of the functions in $\mathcal{F}$. For example, we allow* KeyGen *to take as input a parameter L in such a way that all the parameters in the system depend on L and moreover, the homomorphic evaluation is only supported on circuits (in $\mathcal{F}$) of depth at most L.*

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

$$\Pr\left[f(x) \leftarrow \mathsf{Dec}(\varrho_{\mathsf{sk}}, \mathsf{CT}') \ : \ \substack{(\mathsf{pk},\mathsf{msk},\mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \mathsf{CT} \leftarrow \mathsf{Enc}(\mathsf{pk},x) \\ \mathsf{CT}' \leftarrow \mathsf{Eval}(\mathsf{pk},f,\mathsf{CT})} \right] \geq 1 - \nu(\lambda),$$

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

Correctness of Revocation. Defined as before.

Security. Defined as before (Definition 40).

5.5 Key-Revocable Dual-Regev Encryption

In this section, we present the first construction of key-revocable public-key encryption from standard assumptions. Our construction involves making the Dual Regev public-key encryption of Gentry, Peikert and Vaikuntanathan [68] key revocable.

Construction

We define our Dual-Regev construction below.

Construction 8 (Key-Revocable Dual-Regev Encryption)**.** *Let $n \in \mathbb{N}$ be the security parameter and $m \in \mathbb{N}$. Let $q \geq 2$ be a prime and let $\alpha, \beta, \sigma > 0$ be parameters. The key-revocable public-key scheme* $\mathsf{RevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ *consists of the following* QPT *algorithms:*

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

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

  for some $\mathbf{y} \in \mathbb{Z}_q^n$. *Output* $\mathsf{pk} = (\mathbf{A}, \mathbf{y})$, $|\mathsf{sk}\rangle = |\psi_{\mathbf{y}}\rangle$ *and* $\mathsf{msk} = \mathsf{td}_{\mathbf{A}}$.

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

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

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

- $\mathsf{Revoke}(\mathsf{msk}, \mathsf{pk}, \varrho) \rightarrow \{\top, \bot\}$*: on input* $\mathsf{td}_{\mathbf{A}} \leftarrow \mathsf{msk}$ *and* $(\mathbf{A}, \mathbf{y}) \leftarrow \mathsf{pk}$, *apply the measurement* $\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, \mathbb{1} - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\}$ *onto the state* ϱ *using the procedure* $\mathsf{QSampGauss}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ *in Algorithm 4. Output* $\top$, *if the measurement is successful, and output* $\bot$ *otherwise.*

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

Correctness of Revocation.

Correctness of Revocation. This follows from Theorem 16.

Let us now prove the security of our key-revocable Dual-Regev scheme in Construction 8. Our first result concerns $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda))$-security, i.e., we assume that revocation succeeds with overwhelming probability.

Theorem 24. *Let* $n \in \mathbb{N}$ *and* q *be a prime modulus with* $q = 2^{o(n)}$ *and* $m \geq 2n\log q$, *each parameterized by the security parameter* $\lambda \in \mathbb{N}$. *Let* $\sqrt{8m} < \sigma < q/\sqrt{8m}$ *and let* $\alpha, \beta \in (0, 1)$ *be noise ratios chosen such that* $\beta/\alpha = 2^{o(n)}$ *and* $1/\alpha = 2^{o(n)}\cdot\sigma$. *Then, assuming the subexponential hardness of the* $\mathsf{LWE}_{n,q,\alpha q}^m$ *and* $\mathsf{SIS}_{n,q,\sigma\sqrt{2m}}^m$ *problems, the scheme* $\mathsf{RevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ *in Construction 8 is a* $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda))$*-secure key-revocable public-key encryption scheme according to Definition 40.*

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

Theorem 25. *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and let $\alpha, \beta \in (0,1)$ be noise ratios chosen such that $\beta/\alpha = 2^{o(n)}$ and $1/\alpha = 2^{o(n)} \cdot \sigma$. Assuming Construction 1, the scheme* $\mathsf{RevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ *in Construction 8 is a* $(\mathsf{negl}(\lambda), 1 - 1/\mathsf{poly}(\lambda))$*-secure key-revocable public-key encryption scheme according to Definition 40.*

Guide for proving Theorem 24 and Theorem 25.

- First, we prove a technical lemma (Lemma 28) that helps us remove the condition that revocation succeeds when analyzing the advantage of a distinguisher. Our proof uses *projective implementations* which allow us to estimate the success probability of quantum programs.

- The next step towards proving Theorem 24 is a search-to-decision reduction with quantum auxiliary input for the Dual-Regev scheme (Theorem 30). Here, we show how to extract a short vector mapping $\mathbf{A}$ to $\mathbf{y}$ from an efficient adversary who has a non-negligible distinguishing advantage at distinguishing Dual-Regev ciphertexts from uniform.

- Next, we state the *Simultaneous Dual-Regev Extraction* conjecture in Conjecture 1, which is a strengthening of our search-to-decision reduction in Theorem 30. Informally, it says that extraction of a short vector mapping $\mathbf{A}$ to $\mathbf{y}$ succeeds, even if we apply revocation on a separate register. We prove that Conjecture 1 holds assuming $\mathsf{LWE}/\mathsf{SIS}$ in the special case when revocation succeeds with overwhelming probability. This is captured by Theorem 31.

- Next, we prove technical lemma which exploits the search-to-reduction to extract two *distinct* short vectors mapping $\mathbf{A}$ to $\mathbf{y}$. This is proven in Section 5.5.

- Finally, we put all the pieces together in Section 5.5 and show how to use the result from Section 5.5 in order to break the SIS assumption.

Threshold Implementations

In this section, we prove Lemma 28. This is a key ingredient in the proof of our main theorem, i.e., our simultaneous search-to-decision reduction with quantum auxiliary information in Conjecture 1.

First, we review some recent techniques that allow us to measure the success probability of *quantum programs*. In the classical setting, this task is fairly straightforward: simply execute a given program on samples from a *test distribution*, and check how many times the program succeeds.

Using standard concentration inequalities, one can then estimate the success probability to inverse polymonial precision. In the quantum realm, however, this task is non-trivial if the quantum program is run with respect to quantum auxiliary inputs.

Inspired by the work of Marriott and Watrous [100], Zhandry [138] introduced the notion of projective implementations which allow us to accomplish this task efficiently. Below, we introduce some relevant definitions and results from the original work of Zhandry [138], as well as subsequent follow-up works [6, 54, 18]. First, we discuss *inefficient* measurement techniques for measuring the success probability of a quantum program. Next, we move onto *efficient* measurement techniques that allow us to obtain such estimates approximately.

Inefficient measurements. Suppose we we have quantum program, say consisting of a quantum circuit and some quantum auxiliary input, and we wish to estimate its success probability. A natural starting point is to consider a two-outcome POVM $\mathcal{P} = (P, Q)$ over the two outcomes 0 (success) and 1 (failure). Zhandry [138] showed that for any such $\mathcal{P}$, there exists a natural projective measurement (called a *projective implementation*) such that the post-measurement state corresponds precisely to an eigenvector of P. Moreover, there exists a projective measurement $\mathcal{E}$ that *measures* the success probability with respect to $\mathcal{P}$ on some auxiliary input state; specifically,

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

The measurement $\mathcal{E}$ is projective in the following sense: whenever we apply the same measurement $\mathcal{E}$ on the post-measurement state, we obtain precisely the same outcome. The following theorem is implicit in [138, Lemma 1], but we rely on the presentation from [18, Theorem 2.5].

Theorem 26 (Projective implementation)**.** *Let $\mathcal{P} = (P, Q)$ be a two-outcome POVM and let $\mathcal{D}$ be the distribution over the eigenvalues of P. Then, there exists a projective measurement $\mathcal{E} = \{E_p\}_{p \in \mathcal{D}}$ with index set $\mathcal{D}$ such that: for every quantum state ϱ, where we let $\varrho_p = E_p \varrho E_p$ denote the sub-normalized post-measurement state after measuring ϱ via E_p, it holds that*

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

Remark 27. *Suppose that $\mathcal{P} = (P, Q)$ is a two-outcome POVM and that P has an eigenbasis $\{|\psi_i\rangle\}$ with associated eigenvalues $\{\lambda_i\}$. Because P and Q commute, they share a common eigenbasis. In* *this case, there exists a natural measurement $\mathcal{E}$ that corresponds to a projective implementation of the POVM $\mathcal{P}$; namely, for any input $|\psi\rangle$, which can express as $|\psi\rangle = \sum_i \alpha_i |\psi_i\rangle$, the measurement $\mathcal{E} = \{E_{\lambda_i}\}$ will result in outcome λ_i and a leftover eigenstate $|\psi_i\rangle$ with probabiliy $|\alpha_i|^2$.*

Next, we use a generalization of projective implementations introduced in [6]. Rather than estimating the success probability directly, we can instead measure whether it is above or below a certain threshold. This gives rise to the following notion of *threshold implementations*.

Theorem 28 (Threshold implementation)**.** *Let $\gamma \in (0,1)$ be a parameter and let $\mathcal{P} = (P, Q)$ be a two-outcome POVM, where P has an eigenbasis $\{|\psi_i\rangle\}$ with associated eigenvalues $\{\lambda_i\}$. Then, there exists a projective threshold implementation $(\mathsf{TI}_\gamma(\mathcal{P}), \mathbb{1} - \mathsf{TI}_\gamma(\mathcal{P}))$ such that*

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

The proof of the theorem above follows directly from Theorem 26 by considering the projective measurements $\mathsf{TI}_\gamma(\mathcal{P}) = \sum_{i:\lambda_i \leq \gamma} E_{\lambda_i}$ and $\mathbb{1} - \mathsf{TI}_\gamma(\mathcal{P}) = \mathbb{1} - \sum_{i:\lambda_i > \gamma} E_{\lambda_i}$.

Finally, we also use the following *symmetric* variant of threshold implementations which were considered in [18, Theorem 2.6]. Here, the projective measurement determines whether the success probability is either close to $1/2$ or far from $1/2$.

Theorem 29 (Symmetric threshold implementation)**.** *Let $\gamma \in (0,1/2)$ be a parameter and let $\mathcal{P} = (P, Q)$ be a two-outcome POVM, where P has an eigenbasis $\{|\psi_i\rangle\}$ with associated eigenvalues $\{\lambda_i\}$. Then, there exists a projective threshold implementation $(\mathsf{STI}_\gamma(\mathcal{P}), \mathbb{1} - \mathsf{STI}_\gamma(\mathcal{P}))$ such that*

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

The proof of the theorem above follows directly from Theorem 26 by considering the projective measurements $\mathsf{STI}_\gamma(\mathcal{P}) = \sum_{i:|\lambda_i - \frac{1}{2}| \leq \gamma} E_{\lambda_i}$ and $\mathbb{1} - \mathsf{STI}_\gamma(\mathcal{P}) = \mathbb{1} - \sum_{i:|\lambda_i - \frac{1}{2}| > \gamma} E_{\lambda_i}$.

Efficient measurements. The quantum measurements we described above can, in general, not be implemented efficiently. However, Zhandry [138] showed that there exist so-called efficient *approximate implementations* which allow one to obtain approximate estimates of the success probability of a quantum program. In this section, we review some basic definitions and results that allow us to perform such measurmenents efficiently.

Definition 41 (Mixture of projective measurements)**.** *Let* $\mathcal{P} = \{\mathcal{P}_i\}_{i\in\mathcal{I}}$ *be a collection of binary outcome projective measurements* $\mathcal{P}_i = (P_i, Q_i)$ *over the same Hilbert space* $\mathcal{H}$*, and suppose that* P_i *corresponds to outcome* 1 *and* Q_i *corresponds to outcome* 0*. Let* D *be a distribution over the the index set* $\mathcal{I}$*. Then,* $\mathcal{P}_D = (P_D, Q_D)$ *is the following mixture of pojective measurements:*

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

The following result is adapted from [138, Theorem 6.2] and [6, Corollary 1].

Lemma 26 (Approximate threshold implementation)**.** *Let* $\mathcal{P}_D = (P_D, Q_D)$ *be a binary outcome POVM over Hilbert space* $\mathcal{H}$ *that is a mixture of projective measurements over some distribution* D*. Let* $\varepsilon, \delta, \gamma \in (0,1)$*. Then, there exists an efficient binary-outcome quantum algorithm* $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma}$*, interpreted as the POVM element corresponding to outcome* 1*, such that the following holds:*

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

Finally, we use the following *symmetric* version of the approximate threshold implementation Lemma 26 which is a variant of [18, Theorem 2.8].

Lemma 27 (Symmetric approximate threshold implementation)**.** *Let* $\mathcal{P}_D = (P_D, Q_D)$ *be a binary outcome POVM over Hilbert space* $\mathcal{H}$ *that is a mixture of projective measurements over some distribution* D*. Let* $\gamma \in (0, 1/2)$ *and* $\varepsilon \in (0, \gamma/2)$*, and let* $\delta \in (0,1)$*. Let* D *be a distribution. Then, there exists an efficient binary-outcome quantum algorithm* $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma}$*, interpreted as the POVM element corresponding to outcome* 1*, such that the following holds:*

- *For all quantum states* ϱ*,* $\mathrm{Tr}[\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},D,\gamma-\epsilon}\, \varrho] \geq \mathrm{Tr}[\mathsf{STI}_\gamma(\mathcal{P}_D)\, \varrho] - \delta$*.*

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

Proof of Lemma 28. We are now ready to prove the main result of this subsection.

Lemma 28. *Let* $\lambda \in \mathbb{N}$ *be a parameter and let* $\varrho_{R,\textsc{Aux}}$ *be a quantum state on systems* R *and* $\textsc{Aux}$ *of at most* $\mathsf{poly}(\lambda)$ *many qubits. Let* D_0, D_1 *be two efficiently samplable distributions with support* $\mathcal{X}$*. Let* $\mathcal{D}$ *be a* QPT *algorithm. Suppose that the following two properties hold:*

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

Then, there exists a QPT *algorithm* $\tilde{\mathcal{D}}$ *and a polynomial* $\mu(\lambda)$ *such that* $\tilde{\mathcal{D}}$ *succeeds at distinguishing* D_0 *and* D_1 *with advantage at least* $1/\mu(\lambda)$ *on the reduced system alone, i.e.,*

$$\left|\Pr\left[\tilde{\mathcal{D}}(x, \textsc{Aux}) = b \;:\; \substack{b \xleftarrow{\$} \{0,1\} \\ x \sim D_b}\right] - \frac{1}{2}\right| \geq \frac{1}{\mu(\lambda)},$$

where system $\textsc{Aux}$ *corresponds to the reduced state* $\varrho_{\textsc{Aux}} = \mathrm{Tr}_R[\varrho_{R,\textsc{Aux}}]$*.*

Proof. Consider the binary outcome POVM $\mathcal{P} = (P_{(D_0,D_1)}, Q_{(D_0,D_1)})$ with $Q_{(D_0,D_1)} = \mathbb{1} - P_{(D_0,D_1)}$ which is the following mixture of projective measurements such that

$$P_{(D_0,D_1)} = \frac{\Pi_0 + \Pi_1}{2}$$

where Π_0, Π_1 are mixtures of two-outcome POVMs $\{\mathcal{P}_x\}$ that correspond to running $\mathcal{D}$ on samples x from D_0, D_1 and system $\textsc{Aux}$, and then measuring whether the output is 0 or 1, i.e.,

$$\Pi_0 = \sum_{x \in \mathcal{X}} \Pr[x \leftarrow D_0]\, \mathcal{P}_x \qquad \text{and} \qquad \Pi_1 = \sum_{x \in \mathcal{X}} \Pr[x \leftarrow D_1]\, \mathcal{P}_x.$$

Let $\mathcal{P}$ have an eigenbasis $\{|\psi_i\rangle\}$ with eigenvalues $\{\lambda_i\}$. Without loss of generality we can assume that $\varrho_{R,\textsc{Aux}}$ is a pure state $|\psi\rangle = \sum_i \alpha_i |\psi\rangle$ in systems R and Aux. Moreover, we can write $|\psi\rangle$ as

$$|\psi\rangle = \sum_{i:\, |\lambda_i - \frac{1}{2}| \geq \frac{1}{q}} \alpha_i |\psi_i\rangle + \sum_{i:\, |\lambda_i - \frac{1}{2}| < \frac{1}{q}} \alpha_i |\psi_i\rangle .$$

Let $\varepsilon = 1/8p$, $\delta = 2^{-\lambda}$ and $\gamma = 1/2p$ be parameters. Consider the following distinguisher $\tilde{\mathcal{D}}$:

- Run the efficient approximate threshold implementation $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}$ from Lemma 27 on system Aux for the binary-outcome POVM given by $\mathcal{P}$.
- If the outcome of $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P},(D_0,D_1),\gamma}$ is 1, then run $\mathcal{D}$ on the post-measurement system Aux, and output whatever $\mathcal{D}$ outputs. Otherwise, output a random bit.

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

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

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

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

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

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

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

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

Simultaneous Search-to-Decision Reduction with Quantum Auxiliary Input

Our first result concerns distinguishers with quantum auxiliary input that can distinguish between Dual-Regev samples and uniformly random samples with high probability. In Theorem 30, we give a search-to-decision reduction: we show that such distinguishers can be converted into a quantum extractor that can obtain a Dual-Regev secret key with overwhelming probability. We then state a strenghtening of this extraction property (which we call *Simultaneous Dual-Regev Extraction*) in Conjecture 1. Informally, this property states that extraction is possible even if additionally require that a *revocation* procedure succeeds on a separate register.

While we do not know how to prove Construction 1 under standard assumptions, we prove that *Simultaneous Dual-Regev Extraction* holds assuming LWE/SIS in the special case when revocation succeeds with overwhelming probability. This is captured by Theorem 31.

We first show the following result.

Theorem 30 (Search-to-Decision Reduction with Quantum Auxiliary Input)**.** *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and let $m \geq 2n \log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and let $\alpha, \beta \in (0,1)$ be noise ratios with $\beta/\alpha = 2^{o(n)}$ and $1/\alpha = 2^{o(n)} \cdot \sigma$. Let $\mathcal{A} = \{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}, \nu_\lambda)\}_{\lambda \in \mathbb{N}}$ be any non-uniform quantum algorithm consisting of a family of polynomial-sized quantum circuits*

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

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

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

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

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

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

- If $b = 0$: output $\mathsf{lwe}.\mathsf{Dist}^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)$ defined in Figure 5.3.
- If $b = 1$: output $\mathsf{unif}.\mathsf{Dist}^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)$ defined in Figure 5.4.

Figure 5.2: The experiment $\mathsf{SearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}, b\right)$.

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

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

H_0: This is the distribution $\mathsf{lwe}.\mathsf{Dist}^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)$ in Figure 5.3.

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

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

Figure 5.3: The distribution $\mathsf{lwe}.\mathsf{Dist}^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)$.

H_1: This is the following distribution:

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

5. Run the distinguisher $\mathcal{D}(\mathbf{A}, \mathbf{y}, \mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal, \mathbf{s}^\intercal \mathbf{y} + e', \varrho_{\text{AUX}})$ on the reduced state ϱ_{AUX}.

H_2 : This is the following distribution:

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

H_3 : This is the following distribution:

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

H_4: This is the following distribution:

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

H_5: This is the distribution $\mathsf{u}nif.\mathsf{D}ist^{\mathcal{A}, \mathcal{D}}(1^\lambda)$ in Figure 5.4.

We now show the following:

Claim 10. *Assuming $\mathsf{LWE}^m_{n,q,\alpha q}$, the hybrids H_0 and H_1 are computationally indistinguishable,*

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

$\underline{\mathsf{u}nif.\mathsf{D}ist^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)}$:

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

Figure 5.4: The distribution $\mathsf{u}nif.\mathsf{D}ist^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)$.

Proof. Here, we invoke the *Gaussian-collapsing property* in Corollary 3 which states that the following samples are indistinguishable under $\mathsf{LWE}^m_{n,q,\alpha q}$,

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

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

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

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

Claim 11. *Hybrids* H_1 *and* H_2 *are statistically indistinguishable. In other words,*

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

Proof. Here, we invoke the *noise flooding* property in Lemma 15 to argue that $\mathbf{e}^\intercal\mathbf{x}_0 \ll e'$ holds with overwhelming probability for our choice of parameters. Therefore, the distributions in H_1 and H_2 are computationally indistinguishable. □

Claim 12. *Assuming* $\mathsf{LWE}^m_{n,q,\alpha q}$*, the hybrids* H_2 *and* H_3 *are computationally indistinguishable,*

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

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

Finally, we show the following:

Claim 13. *Assuming* $\mathsf{LWE}^m_{n,q,\alpha q}$*, the hybrids* H_4 *and* H_5 *are computationally indistinguishable,*

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

Proof. Here, we invoke the *Gaussian-collapsing property* in Corollary 3 again. □

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

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

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

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

Procedure:

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

Figure 5.5: The distinguisher $\tilde{\mathcal{D}}(\mathbf{A}, \mathbf{y}, \mathbf{u}, v, \varrho)$.

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

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

From Theorem 20, it follows that there exists a Goldreich-Levin extractor $\mathcal{E}$ running in time $T(\mathcal{E}) = \mathsf{poly}(\lambda, n, m, \sigma, q, 1/\varepsilon)$ that outputs a short vector in $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ with probability at least

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

Assuming the $\mathsf{LWE}_{n,q,\alpha q}^m$ assumption, we can invoke the Gaussian-collapsing property in Corollary 3 once again which implies that the quantum extractor $\mathcal{E}$ satisfies

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

This proves the claim. □

Next, we give a strengthening of our result in Theorem 30 and state a *simultaneous* search-to-decision reduction with quantum auxiliary input which holds even if additionally require that a *revocation* procedure succeeds on a separate register.

To formalize the notion that revocation is applied on a separate register, we introduce the following procedure called IneffRevoke which is defined below.

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

Conjecture 1 (Simultaneous Dual-Regev Extraction (SDRE)). *Let $n \in \mathbb{N}$ be the security parameter. There exist parameters (each parameterized by λ), where q is a prime modulus with $q = 2^{o(n)}$, $m \geq 2n \log q$, $\sqrt{8m} < \sigma < q/\sqrt{8m}$, $\alpha, \beta \in (0,1)$ with $\beta/\alpha = 2^{o(n)}$, and $1/\alpha = 2^{o(n)} \cdot \sigma$, such that the following holds: Let $\mathcal{A} = \{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}, \nu_\lambda)\}_{\lambda \in \mathbb{N}}$ be any non-uniform quantum algorithm consisting of a family of polynomial-sized quantum circuits*

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

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

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

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

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

Procedure:

1. Apply the (inefficient) projective measurement

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

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

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

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

Figure 5.6: The procedure $\mathsf{IneffRevoke}(\mathbf{A}, \mathbf{y}, \sigma, \varrho_{\mathrm{R}})$.

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

- If $b = 0$: output $\mathsf{s}imult.\mathsf{l}we.\mathsf{D}ist^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 5.8.
- If $b = 1$: output $\mathsf{s}imult.\mathsf{u}nif.\mathsf{D}ist^{\mathcal{A},\mathcal{D}}\left(1^\lambda\right)$ defined in Figure 5.9.

Figure 5.7: The experiment $\mathsf{SimultSearchToDecisionExpt}^{\mathcal{A},\mathcal{D}}\left(1^\lambda, b\right)$.

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

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

Towards a proof of the conjecture. We now give a proof of *Simultaneous Dual-Regev Extraction* (Construction 1) in the special case when revocation succeeds with overwhelming probability.

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

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

Figure 5.8: The distribution $\mathsf{s}imult.\mathsf{lwe}.\mathsf{D}ist^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)$.

```
$\underline{\mathsf{s}imult.\mathsf{u}nif.\mathsf{D}ist^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right):}$

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

Figure 5.9: The distribution $\mathsf{s}imult.\mathsf{u}nif.\mathsf{D}ist^{\mathcal{A},\mathcal{D}}\left(1^{\lambda}\right)$.

Theorem 31. *Let $n \in \mathbb{N}$. Let q be a prime with $q = 2^{o(n)}$, $m \geq 2n\log q$, $\sqrt{8m} < \sigma < q/\sqrt{8m}$, and let $\alpha, \beta \in (0,1)$ with $\beta/\alpha = 2^{o(n)}$ with $1/\alpha = 2^{o(n)} \cdot \sigma$. Let $\mathcal{A} = \{(\mathcal{A}_{\lambda,\mathbf{A},\mathbf{y}}, \nu_{\lambda})\}_{\lambda\in\mathbb{N}}$ be any* *non-uniform quantum algorithm consisting of a family of polynomial-sized quantum circuits*

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

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

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

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

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

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

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

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

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

We can now invoke Lemma 28 to argue that there exists a QPT distinguisher $\tilde{\mathcal{D}}$ (that internally runs $\mathcal{D}$) and succeeds on the reduced system Aux alone, i.e.

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

for some $\bar{\varepsilon} = 1/\mathsf{poly}(\lambda)$. In other words, the QPT algorithm $\tilde{\mathcal{D}}$ can successfully predict whether it has received a Dual-Regev sample or a uniformly random sample. Therefore, we can now invoke Theorem 30 to argue there exists an extractor $\mathcal{E}$ that takes as input $\mathbf{A}$, $\mathbf{y}$ and system $\textsc{Aux}$ of the state $\varrho_{\mathrm{R},\textsc{Aux}}$ and outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\mathsf{poly}(\lambda, m, \sigma, q, 1/\bar{\varepsilon})$ such that

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

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

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

Using Bonferroni's inequality, we can argue that

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

This proves the claim.

□

Distinct Pair Extraction

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

Lemma 29 (Distinct pair extraction). *Let $\varrho \in \mathcal{D}(\mathcal{H}_X \otimes \mathcal{H}_Y)$ be an any density matrix, for some Hilbert spaces $\mathcal{H}_X$ and $\mathcal{H}_Y$. Let $|\psi\rangle = \sum_{x\in\mathcal{S}} \alpha_x |x\rangle \in \mathcal{H}_X$ be any state supported on a subset $\mathcal{S} \subseteq \mathcal{X}$, and let $\mathbf{\Pi} = |\psi\rangle\langle\psi|$ denote its associated projection. Let $\mathbf{\Pi}_{\mathcal{S}}$ be the projector onto $\mathcal{S}$ with*

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

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

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

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

$$\mathbf{\Gamma} = \sum_{x,x'\in\mathcal{S}:x\neq x'} |x\rangle\langle x|_X \otimes V^\dagger_{Y\to X'E}(|x'\rangle\langle x'|_{X'} \otimes \mathbb{1}_E)V_{Y\to X'E}.$$

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

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

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

Proof. Because the order in which we apply $\mathbf{\Gamma}$ and $(\mathbf{\Pi}\otimes\mathbb{1})$ does not matter, we have the inequality

$$\mathrm{Tr}\left[\mathbf{\Gamma}\varrho\right] \geq \mathrm{Tr}\left[(\mathbf{\Pi}\otimes\mathbb{1})\,\mathbf{\Gamma}\varrho\right] = \mathrm{Tr}\left[(\mathbf{\Pi}\otimes\mathbb{1})\,\mathbf{\Gamma}\varrho\,(\mathbf{\Pi}\otimes\mathbb{1})\right] = \mathrm{Tr}\left[\mathbf{\Gamma}(\mathbf{\Pi}\otimes\mathbb{1})\varrho\,(\mathbf{\Pi}\otimes\mathbb{1})\right]. \tag{5.1}$$

Notice also that $(\mathbf{\Pi}\otimes\mathbb{1})\varrho(\mathbf{\Pi}\otimes\mathbb{1})$ lies in the image of $(\mathbf{\Pi}\otimes\mathbb{1})$ with $\mathbf{\Pi} = |\psi\rangle\langle\psi|$, and thus

$$(\mathbf{\Pi}\otimes\mathbb{1})\varrho(\mathbf{\Pi}\otimes\mathbb{1}) = \mathrm{Tr}[(\mathbf{\Pi}\otimes\mathbb{1})\varrho] \cdot (|\psi\rangle\langle\psi| \otimes \sigma), \tag{5.2}$$

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

$$\begin{aligned}
\mathrm{Tr}\left[\mathbf{\Gamma}\varrho\right] &\geq \mathrm{Tr}\left[\mathbf{\Gamma}(\mathbf{\Pi}\otimes\mathbb{1})\varrho\,(\mathbf{\Pi}\otimes\mathbb{1})\right] && \text{(using inequality (5.1))}\\
&= \mathrm{Tr}[(\mathbf{\Pi}\otimes\mathbb{1})\varrho] \cdot \mathrm{Tr}\left[\mathbf{\Gamma}\,(|\psi\rangle\langle\psi| \otimes \sigma)\right] && \text{(using equation (5.2))}\\
&= \mathrm{Tr}[\mathbf{\Pi}\varrho_X] \cdot \mathrm{Tr}\left[\sum_{x,x'\in\mathcal{S}:x\neq x'} |x\rangle\langle x|_X \otimes V^\dagger_{Y\to X'E}\,(|x'\rangle\langle x'|_{X'} \otimes \mathbb{1}_E)\,V_{Y\to X'E}\,(|\psi\rangle\langle\psi| \otimes \sigma)\right]\\
&= \mathrm{Tr}[\mathbf{\Pi}\varrho_X] \cdot \sum_{x'\in\mathcal{S}}\left(\sum_{x\in\mathcal{S}:x\neq x'} |\langle x|\psi\rangle|^2\right)\mathrm{Tr}\left[V^\dagger_{Y\to X'E}(|x'\rangle\langle x'|_{X'} \otimes \mathbb{1}_E)V_{Y\to X'E}\,\sigma\right]\\
&= \mathrm{Tr}[\mathbf{\Pi}\varrho_X] \cdot \sum_{x'\in\mathcal{S}}\left(1 - |\alpha_{x'}|^2\right)\mathrm{Tr}\left[(|x'\rangle\langle x'|_{X'} \otimes \mathbb{1}_E)V_{Y\to X'E}\,\sigma\,V^\dagger_{Y\to X'E}\right]\\
&\geq \mathrm{Tr}[\mathbf{\Pi}\varrho_X] \cdot \left(1 - \max_{x\in\mathcal{S}}|\alpha_x|^2\right) \cdot \sum_{x'\in\mathcal{S}}\mathrm{Tr}\left[(|x'\rangle\langle x'|_{X'} \otimes \mathbb{1}_E)V_{Y\to X'E}\,\sigma\,V^\dagger_{Y\to X'E}\right]\\
&= \mathrm{Tr}[\mathbf{\Pi}\varrho_X] \cdot \left(1 - \max_{x\in\mathcal{S}}|\alpha_x|^2\right) \cdot \sum_{x'\in\mathcal{S}}\mathrm{Tr}\left[|x'\rangle\langle x'|_{X'}\,\mathrm{Tr}_E\left[V_{Y\to X'E}\,\sigma\,V^\dagger_{Y\to X'E}\right]\right]\\
&= \mathrm{Tr}[\mathbf{\Pi}\varrho_X] \cdot \left(1 - \max_{x\in\mathcal{S}}|\alpha_x|^2\right) \cdot \mathrm{Tr}\left[\mathbf{\Pi}_{\mathcal{S}}\,\mathcal{E}_{Y\to X'}(\sigma)\right].
\end{aligned}$$

This proves the claim. □

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

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

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

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

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

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

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

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

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

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

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

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

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

Figure 5.10: The key-revocable security experiment according to Definition 40.

Proof of Theorem 24

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

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

for some $\varepsilon(\lambda)$ with respect to $\mathsf{Expt}_{\mathcal{A}}(1^\lambda, b)$ in Figure 5.10. We show that $\varepsilon(\lambda)$ is negligible.

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

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

We show that we can use $\mathcal{A}$ to break the $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ problem. Without loss of generality, we assume that $\mathcal{A}$ submits the plaintext $x = 0$. By the assumption that revocation succeeds with overwhelming probability and since $\epsilon(\lambda) \geq 1/\mathsf{poly}(\lambda)$, we can use Theorem 31 to argue that there exists a quantum Goldreich-Levin extractor $\mathcal{E}$ that takes as input $\mathbf{A}$, $\mathbf{y}$ and system $\textsc{Aux}$ of the state $\varrho_{R,\textsc{Aux}}$ and outputs a short vector in the coset $\Lambda_q^{\mathbf{y}}(\mathbf{A})$ in time $\mathsf{poly}(\lambda, m, \sigma, q, 1/\varepsilon)$ such that

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

Here, we rely on the correctness of $\mathsf{GenTrap}$ in Theorem 2 and $\mathsf{QSampGauss}$ in Theorem 16.

Consider the following procedure in Algorithm 5.

```
Algorithm 5: SIS_Solver($\mathbf{A}$)

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

Generate a Gaussian state $(|\psi_{\mathbf{y}}\rangle, \mathbf{y}) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ with
    $|\psi_{\mathbf{y}}\rangle = \sum_{\mathbf{x} \in \mathbb{Z}_q^m,\ \mathbf{A}\mathbf{x} = \mathbf{y} \pmod{q}} \varrho_\sigma(\mathbf{x}) \ |\mathbf{x}\rangle$
  for some vector $\mathbf{y} \in \mathbb{Z}_q^n$.
Run $\mathcal{A}$ to generate a bipartite state $\varrho_{R\,\textsc{Aux}}$ in systems $\mathcal{H}_{\mathrm{R}} \otimes \mathcal{H}_{\textsc{Aux}}$ with $\mathcal{H}_{\mathrm{R}} = \mathcal{H}_q^m$.
Measure system R in the computational basis, and let $\mathbf{x}_0 \in \mathbb{Z}_q^n$ denote the outcome.
Run the quantum Goldreich-Levin extractor $\mathcal{E}(\mathbf{A}, \mathbf{y}, \varrho_{\textsc{Aux}})$ from Conjecture 1, where $\varrho_{\textsc{Aux}}$ is
  the reduced state in system $\mathcal{H}_{\textsc{Aux}}$, and let $\mathbf{x}_1 \in \mathbb{Z}_q^n$ denote the outcome.
Output the vector $\mathbf{x} = \mathbf{x}_1 - \mathbf{x}_0$.
```

To conclude the proof, we show that $\mathsf{SIS\_Solver}(\mathbf{A})$ in Algorithm 5 breaks the $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ problem whenever $\varepsilon(\lambda) = 1/\mathsf{poly}(\lambda)$. In order to guarantee that $\mathsf{SIS\_Solver}(\mathbf{A})$ is successful, we use the distinct pair extraction result of Lemma 29. This allows us to analyze the probability of simultaneously extracting two distinct short pre-images $\mathbf{x}_0 \neq \mathbf{x}_1$ such that $\mathbf{A}\mathbf{x}_0 = \mathbf{y} = \mathbf{A}\mathbf{x}_1 \pmod q$ – both in terms of the success probability of revocation and the success probability of extracting a pre-image from the adversary's state ϱ_{Aux} in system $\mathcal{H}_{\text{Aux}}$. Assuming that $\mathbf{x}_0, \mathbf{x}_1$ are distinct short pre-images such that $\|\mathbf{x}_0\| \leq \sigma\sqrt{\frac{m}{2}}$ and $\|\mathbf{x}_1\| \leq \sigma\sqrt{\frac{m}{2}}$, it then follows that the vector $\mathbf{x} = \mathbf{x}_1 - \mathbf{x}_0$ output by $\mathsf{SIS\_Solver}(\mathbf{A})$ has norm at most $\sigma\sqrt{2m}$, and thus yields a solution to $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$.

We remark that the state $|\psi_{\mathbf{y}}\rangle$ prepared by Algorithm 5 is not normalized for ease of notation. Note that the tail bound in Lemma 11 implies that (the normalized variant of) $|\psi_{\mathbf{y}}\rangle$ is within negligible trace distance of the state with support $\{\mathbf{x} \in \mathbb{Z}_q^m : \|\mathbf{x}\| \leq \sigma\sqrt{\frac{m}{2}}\}$. Therefore, for the sake of Lemma 29, we can assume that $|\psi_{\mathbf{y}}\rangle$ is a normalized state of the form

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

Before we analyze Algorithm 5, we first make two technical remarks. First, since $\sigma \geq \omega(\sqrt{\log m})$, it follows from Lemma 17 that, for any full-rank $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and for any $\mathbf{y} \in \mathbb{Z}_q^n$, we have

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

Second, we can replace the procedure $\mathsf{Revoke}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \varrho_R)$ by an (inefficient) projective measurement $\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, I - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\}$, since they produce statistically close outcomes. This follows from the fact that $\mathsf{Revoke}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \varrho_{\mathrm{R}})$ applies the procedure $\mathsf{QSampGauss}$ in Algorithm 4 as a subroutine, which is correct with overwhelming probability acccording to Theorem 16.

Let us now analyze the success probability of Algorithm 5. Putting everything together, we get

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

In the last line, we applied the simultaneous search-to-decision reduction from Theorem 31 and Lemma 8. Therefore, $\mathsf{SIS\_Solver}(\mathbf{A})$ in Algorithm 5 runs in time $\mathsf{poly}(q,1/\varepsilon)$ and solves $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$ whenever $\varepsilon=1/\mathsf{poly}(\lambda)$. Therefore, we conclude that $\varepsilon(\lambda)$ must be negligible. □

Proof of Theorem 25

Proof. The proof is the same as in Theorem 24, except that we invoke Conjecture 1 instead of Theorem 31 to argue that simultaneous extraction succeeds with sufficiently high probability. □

5.6 Key-Revocable Fully Homomorphic Encryption

In this section, we describe our key-revocable (leveled) fully homomorphic encryption scheme from LWE which is based on the so-called DualGSW scheme used by Mahadev [99] which itself is a variant of the homomorphic encryption scheme by Gentry, Sahai, and Waters [69].

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

Note that the associated (non-linear) inverse operation $\mathbf{G}^{-1}$ converts vectors in $\mathbb{Z}_q^{m+1}$ to their binary representation in $\{0,1\}^N$. In other words, we have that $\mathbf{G} \circ \mathbf{G}^{-1}$ acts as the identity operator.

Construction

Let us now describe our key-revocable fully homomorphic encryption scheme.

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

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

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

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

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

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

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

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

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

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

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

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

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

Figure 5.11: The key-revocable security experiment according to Definition 40.

Construction 9 (Key-Revocable $\mathsf{DualGSW}$ encryption)**.** *Let $\lambda \in \mathbb{N}$ be a parameter. The scheme* $\mathsf{RevDualGSW} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Revoke})$ *consists of the following* QPT *algorithms:*

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

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

for some $\mathbf{y} \in \mathbb{Z}_q^n$*. Output* $\mathsf{pk} = (\mathbf{A}, \mathbf{y})$*,* $|\mathsf{sk}\rangle = |\psi_{\mathbf{y}}\rangle$ *and* $\mathsf{msk} = \mathsf{td}_{\mathbf{A}}$.

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

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

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

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

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

$\mathsf{Revoke}(\mathsf{msk}, \mathsf{pk}, \varrho) \to \{\top, \bot\}$*: on input* $\mathsf{td}_{\mathbf{A}} \leftarrow \mathsf{msk}$ *and* $(\mathbf{A}, \mathbf{y}) \leftarrow \mathsf{pk}$*, apply the projective measurement* $\{|\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|, \mathbb{1} - |\psi_{\mathbf{y}}\rangle\langle\psi_{\mathbf{y}}|\}$ *onto ϱ using* $\mathsf{SampGauss}(\mathbf{A}, \mathsf{td}_{\mathbf{A}}, \mathbf{y}, \sigma)$ *in Algorithm 4. Output $\top$ if the measurement is successful, and output $\bot$ otherwise.*

Proof of security

Our first result on the security of Construction 9 concerns $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda))$-security, i.e., we assume that revocation succeeds with overwhelming probability.

Theorem 32. *Let L be an upper bound on the* NAND*-depth of the circuit which is to be evaluated. Let $n \in \mathbb{N}$ and q be a prime modulus with $n = n(\lambda, L) \gg L$, $q = 2^{o(n)}$ and $m \geq 2n\log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $N = (m+1)\cdot\lceil\log q\rceil$ be an integer. Let $q/\sqrt{8m} > \sigma > \sqrt{8m}$ and let $\alpha, \beta \in (0, 1)$ be parameters such that $\beta/\alpha = 2^{o(n)}$ and $1/\alpha = 2^{o(n)}\cdot\sigma$. Then, assuming the subexponential hardness of the* $\mathsf{LWE}_{n,q,\alpha q}^m$ *and* $\mathsf{SIS}_{n,q,\sigma\sqrt{2m}}^m$ *problems, the scheme* $\mathsf{RevDualGSW} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Revoke})$ *in Construction 9 is a* $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda))$*-secure key-revocable (leveled) fully homomorphic encryption scheme according to Definition 40.*

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

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

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

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

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

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

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

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

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

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

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

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

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

Figure 5.12: The key-revocable security experiment according to Definition 40.

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

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

for some $\varepsilon(\lambda)$ with respect to $\mathsf{Expt}_{\mathcal{A}}(1^\lambda, b)$ in Figure 5.12. Note that the $\mathsf{RevDualGSW}$ ciphertext can (up to an additive shift) be thought of as a column-wise concatenation of N-many independent ciphertexts of our key-revocable Dual-Regev scheme in Construction 8. Therefore, we can invoke Theorem 24 in order to argue that $\varepsilon(\lambda)$ is at most negligible.

□

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

Theorem 33. *Let L be an upper bound on the* NAND*-depth of the circuit which is to be evaluated. Let $n \in \mathbb{N}$ and q be a prime modulus with $n = n(\lambda, L) \gg L$, $q = 2^{o(n)}$ and $m \geq 2n\log q$, each parameterized by the security parameter $\lambda \in \mathbb{N}$. Let $N = (m+1)\cdot\lceil\log q\rceil$ be an integer. Let $q/\sqrt{8m} > \sigma > \sqrt{8m}$ and let $\alpha, \beta \in (0,1)$ be parameters such that $\beta/\alpha = 2^{o(n)}$ and $1/\alpha = 2^{o(n)} \cdot \sigma$. Then, assuming assuming Conjecture 1, the scheme* $\mathsf{RevDualGSW} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval}, \mathsf{Revoke})$ *in Construction 9 is a $(\mathsf{negl}(\lambda), 1-1/\mathsf{poly}(\lambda))$-secure key-revocable (leveled) fully homomorphic encryption scheme according to Definition 40.*

Proof. The proof is the same as in the theorem before, except that we invoke Theorem 25 instead of Theorem 24 in order to argue security. □

5.7 Revocable Pseudorandom Functions

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

Definition

Let us first recall the traditional notion of PRF security [73], defined as follows.

Definition 42 (Pseudorandom Function)**.** *Let $\lambda \in \mathbb{N}$ and $\kappa(\lambda), \ell(\lambda)$ and $\ell'(\lambda)$ be polynomials. A (post-quantum) pseudorandom function* (pqPRF) *is a pair* (Gen, PRF) *of* PPT *algorithms given by*

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

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

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

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

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

Definition 43 (Revocable Pseudorandom Function)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter and let $\kappa(\lambda), \ell(\lambda)$ and $\ell'(\lambda)$ be polynomials. A revocable pseudorandom function* (rPRF) *is a scheme* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *consisting of the following efficient algorithms:*

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

We additionally require that the following holds:

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

- (Correctness of evaluation:)

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

- (Correctness of revocation:)

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

Security
We define revocable PRF security below.

Definition 44 (Revocable PRF Security)**.** *A revocable pseudorandom function* (rPRF) *consisting of a tuple of* QPT *algorithms* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *has $(\varepsilon, \delta, \mu)$ revocable* PRF *security if, for every* QPT *adversary $\mathcal{A}$ with*

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

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

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

where $\mathsf{Expt}_{\mathcal{A},\mu}$ is as defined in Figure 5.13. If $\delta(\lambda) = 1 - 1/\mathsf{poly}(\lambda)$, $\varepsilon(\lambda) = \mathsf{negl}(\lambda)$ we oftentimes drop (δ, ε) and simply refer to it as rPRF *satisfies μ-revocable* PRF *security.*

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

Initialization Phase:

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

Revocation Phase:

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

Guessing Phase:

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

Figure 5.13: Revocable PRF security

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

Claim 14. *Supoose an* rPRF *scheme* (Gen, PRF, Eval, Revoke) *satisfies* 1*-revocable* PRF *security. Then,* rPRF *also satisfies the stronger notion of (multi-query) revocable PRF security.*

Proof. Let $\mathcal{A}$ be a QPT adversary that participating in the revocable PRF security experiment defined in Figure 5.13 and let $(x_1, y_1), \ldots, (x_\mu, y_\mu)$ denote the challenge input-output pairs, for some polynomial $\mu = \mu(\lambda)$. In the following, we denote by k the PRF key sampled using Gen by the challenger in Figure 5.13. We consider a sequence of hybrids defined as follows.

H_i, for $i \in [\mu + 1]$: In this hybrid, $y_1, \ldots, y_{i-1}$ are sampled uniformly at random from $\{0, 1\}^{\ell'}$ and

$y_i, \ldots, y_\mu$ are generated as follows: $y_j = \mathsf{PRF}(k, x_j)$ for $j \geq i$.

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

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

1. Receive the state ϱ_k from the challenger.
2. Sample $x_{i+1}, \ldots, x_\mu$ uniformly at random from $\{0,1\}^\ell$. Denote $\varrho_k^{(i+1)} = \varrho_k$. Do the following for $j = i+1, \ldots, \mu$: $\mathsf{Eval}(\varrho_k^{(j)}, x_j)$ to obtain y_j. Using the "Almost As Good As New" [3], recover $\varrho_k^{(j+1)}$, where $\varrho_k^{(j+1)}$ is negligibly[10] close to ϱ_k in trace distance.
3. Forward the state $\varrho_k^{(\mu+1)}$ to $\mathcal{A}$.
4. When the challenger submits the message `REVOKE`, forward the same message to $\mathcal{A}$.
5. If $\mathcal{A}$ sends σ, then forward the same state σ to the challenger.
6. If the revocation did not fail, the guessing phase begins. The challenger sends (x^*, y^*). Then, sample $x_1, \ldots, x_{i-1}$ uniformly at random from $\{0,1\}^\ell$ and $y_1, \ldots, y_{i-1}$ uniformly at random from $\{0,1\}^{\ell'}$. Set $x_i = x^*$ and $y_i = y^*$. Send $(x_1, y_1), \ldots, (x_\mu, y_\mu)$ to $\mathcal{A}$.
7. Output b, where b is the output of $\mathcal{A}$.

From the quantum union bound (Lemma 2), the "Almost As Good As New" lemma (Lemma 1) and the correctness of rPRF, it follows that $\mathrm{TD}(\varrho_k, \varrho_k^{(\mu+1)}) \leq \mathsf{negl}(\lambda)$ and thus, the advantage of $\mathcal{A}$ when given $\varrho_k^{(\mu+1)}$ instead of ϱ_k is now at least $\varepsilon - \mathsf{negl}(\lambda)$. Moreover, by the design of $\mathcal{B}$, it follows that the success probability of $\mathcal{B}$ in breaking 1-revocation security of rPRF is exactly the same as the success probability of $\mathcal{A}$ in breaking revocation security of rPRF. This contradicts the fact that rPRF satisfies 1-revocation security.

□

Remark 34. *Our notion of revocable* PRF *security from Definition 44 does not directly imply traditional notion of* pqPRF *security*[11] *from Definition 42. The reason is that the definition does not*

[10] Technically, this depends on the correctness error and we start with a rPRF that is correct with probability negligibly close to 1.

[11] Although any revocable PRF is a *weak* PRF. Recall that a weak PRF is one where the adversary receives as input $(x_1, y_1), \ldots, (x_\mu, y_\mu)$, where x_is are picked uniformly at random. The goal of the adversary is to distinguish the two cases: all y_is are pseudorandom or all y_is are picked uniformly at random.

preclude the possibility of there being an input x (say an all zeroes string) on which, PRF *outputs x itself (or the first bit of x if the output of* PRF *is a single bit).*

Motivated by Theorem 34, we now introduce the following notion of a *strong* rPRF.

Definition 45 (Strong rPRF)**.** *We say that a scheme* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *is a strong revocable pseudorandom function (or, strong* rPRF*) if the following two properties hold:*

1. $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *satisfy revocable* PRF *security according to Definition 44, and*
2. $(\mathsf{Gen}, \mathsf{PRF})$ *satisfy* pqPRF *security according to Definition 42.*

Remark 35. *When instantiating pseudorandom functions in the textbook construction of private-key encryption [72] from revocable pseudorandom functions, we immediately obtain a revocable private-key encryption scheme.*

We show that the issue raised in Theorem 34 is not inherent. In fact, we give a simple generic transformation that allows us to obtain strong rPRFs by making use of traditional pqPRFs.

Claim 15 (Generic Transformation for Strong rPRFs)**.** *Let* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *be an* rPRF *scheme which satisfies revocable* PRF *security, and let* $(\overline{\mathsf{Gen}}, \overline{\mathsf{PRF}})$ *be a* pqPRF*. Then, the scheme* $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}}, \widetilde{\mathsf{Eval}}, \widetilde{\mathsf{Revoke}})$ *is a strong* rPRF *which consists of the following algorithms:*

- $\widetilde{\mathsf{Gen}}(1^\lambda)$*: on input the security parameter 1^λ, first run* $(k, \varrho_k, \mathsf{msk}) \leftarrow \mathsf{Gen}(1^\lambda)$ *and then output* $((K, k), (K, \varrho_k), \mathsf{msk})$*, where* $K \leftarrow \overline{\mathsf{Gen}}(1^\lambda)$ *is a* pqPRF *key.*
- $\widetilde{\mathsf{PRF}}((K, k), x)$*: on input a key (K, k) and string $x \in \{0, 1\}^\ell$, output* $\overline{\mathsf{PRF}}(K, x) \oplus \mathsf{PRF}(k, x)$*.*
- $\widetilde{\mathsf{Eval}}((K, \varrho_k), x)$*: on input (K, ϱ_k) and $x \in \{0, 1\}^\ell$, output* $\overline{\mathsf{PRF}}(K, x) \oplus \mathsf{Eval}(\varrho_k, x)$*.*
- $\widetilde{\mathsf{Revoke}}(\mathsf{msk}, (K, \sigma))$*: on input a master secret key* msk *and a pair (K, ϱ_k), first discard the key K and then run* $\mathsf{Revoke}(\mathsf{msk}, \sigma)$*.*

Proof. Let us first show that the scheme $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}}, \widetilde{\mathsf{Eval}}, \widetilde{\mathsf{Revoke}})$ maintains revocable PRF security. Suppose there exists a QPT adversary $\mathcal{A}$ and a polynomial $\mu = \mu(\lambda) \in \mathbb{N}$ such that

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

for some function $\epsilon(\lambda) = 1/\mathsf{poly}(\lambda)$, and where $\mathsf{Expt}_{\mathcal{A},\mu}$ is the experiment from Figure 5.13 with respect to the scheme $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}}, \widetilde{\mathsf{Eval}}, \widetilde{\mathsf{Revoke}})$. We show that this implies the existence of a QPT distinguisher $\mathcal{D}$ that breaks the revocable PRF security of the scheme $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$.

The distinguisher $\mathcal{D}$ proceeds as follows:

1. $\mathcal{D}$ receives as input a quantum state ϱ_k, where $(k, \varrho_k, \mathsf{msk}) \leftarrow \mathsf{Gen}(1^\lambda)$ is generated by the challenger. Then, $\mathcal{D}$ generates a pqPRF key $K \leftarrow \overline{\mathsf{Gen}}(1^\lambda)$ and sends (K, ϱ_k) to $\mathcal{A}$.

2. When $\mathcal{A}$ returns a state ϱ, $\mathcal{D}$ forwards it to the challenger as part of the revocation phase.

3. When $\mathcal{D}$ receives the challenge input $(x_1, \ldots, x_\mu)$ and $(y_1, \ldots, y_\mu)$ from the challenger, $\mathcal{D}$ sends $(x_1, \ldots, x_\mu)$ and $(\overline{\mathsf{PRF}}(K, x_1) \oplus y_1, \ldots, \overline{\mathsf{PRF}}(K, x_\mu) \oplus y_\mu)$ to $\mathcal{A}$.

4. When $\mathcal{A}$ outputs b', so does the distinguisher $\mathcal{D}$.

Note that the simulated challenge distribution above precisely matches the challenge distribution from the experiment $\mathsf{Expt}_{\mathcal{A},\mu}$ from Figure 5.13. Therefore, if $\mathcal{A}$ succeeds with inverse polynomial advantage $\epsilon(\lambda) = 1/\mathsf{poly}(\lambda)$, so does $\mathcal{D}$ – thereby breaking the revocable PRF security of the scheme $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$. Consequently, $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}}, \widetilde{\mathsf{Eval}}, \widetilde{\mathsf{Revoke}})$ satisfies revocable PRF security.

To see why $(\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{PRF}})$ satisfy pqPRF security according to Definition 42, we can follow a similar argument as above to break the pqPRF security of $(\overline{\mathsf{Gen}}, \overline{\mathsf{PRF}})$. Here, we rely on the fact that the keys $(k, \varrho_k, \mathsf{msk}) \leftarrow \mathsf{Gen}(1^\lambda)$ and $K \leftarrow \overline{\mathsf{Gen}}(1^\lambda)$ are sampled independently from another.

□

Remark 36. *Previous works [54, 90] do not explicitly require in their definitions that either secure software leasing or copy-protection of pseudorandom functions must necessarily preserve the pseudorandomness property (although their constructions could still satisfy such a traditional pseudorandomness property).*

Construction

We construct a PRF satisfying 1-revocation security (Definition 44).

Shift-Hiding Construction. We construct a *shift-hiding* function which is loosely inspired by shift-hiding shiftable functions introduced by Peikert and Shiehian [107].

Let $n, m \in \mathbb{N}$, $q \in \mathbb{N}$ be a modulus and let $\ell = nm\lceil \log q \rceil$. In the following, we consider matrix-valued functions $F : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$, where F is one of the following functions:

- $\mathcal{Z} : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$ which, on input $x \in \{0,1\}^\ell$, outputs an all zeroes matrix $\mathbf{0} \in \mathbb{Z}_q^{n \times m}$, or:
- $H_r : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$ which, on input $x \in \{0,1\}^\ell$, outputs $\mathbf{M} \in \mathbb{Z}_q^{n \times m}$, where $r \in \{0,1\}^\ell$ and $x = r \oplus \mathsf{bindecomp}(\mathbf{M})$, where $\mathbf{M} \in \mathbb{Z}_q^{n \times m}$ and $\mathsf{bindecomp}(\cdot)$ takes as input a matrix and outputs a binary string that is obtained by concatenating the binary decompositions of all the elements in the matrix (in some order).

We show that there exist PPT algorithms $(\mathcal{KG}, \mathcal{E})$ (formally defined in Construction 10) with the following properties:

- $\mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$: on input $1^n, 1^m$, a modulus $q \in \mathbb{N}$, a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and a function $F \in \{\mathcal{Z}\} \cup \{H_r : r \in \{0,1\}^\ell\}$, it outputs a pair of keys (pk_F, sk_F).
- $\mathcal{E}(pk_F, x)$: on input pk_F, $x \in \{0,1\}^\ell$, it outputs $\mathbf{S}_x\mathbf{A} + \mathbf{E}_x + F(x)$, where $\mathbf{S}_x \in \mathbb{Z}_q^{n\times n}$ and $\mathbf{E}_x \in \mathbb{Z}_q^{n\times m}$, where $||\mathbf{E}_x|| \leq nm^2\sigma\lceil\log(q)\rceil$. Moreover, there is an efficient algorithm that recovers $\mathbf{S}_x$ given sk_F and x.

We show that $(\mathcal{KG}, \mathcal{E})$ satisfies a *shift-hiding property*; namely, for any $r \in \{0,1\}^\ell$,

$$\{pk_{\mathcal{Z}}\} \approx_c \{pk_{H_r}\},$$

for any pk_F with $(pk_F, sk_F) \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$, where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, and $F \in \{\mathcal{Z}, H_r\}$.

In the construction below, we consider a bijective function $\phi : [n] \times [m] \times [\lceil\log(q)\rceil] \to [\ell]$.

Construction 10. *Consider the* PPT *algorithms* $(\mathcal{KG}, \mathcal{E})$ *defined as follows:*

- $\mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$: *on input* $1^n, 1^m$, *a modulus* $q \in \mathbb{N}$, *a matrix* $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ *and function* $F \in \{\mathcal{Z}\} \cup \{H_r : r \in \{0,1\}^\ell\}$, *it outputs a pair of keys* $\kappa_F = (pk_F, sk_F)$ *generated as follows:*
  1. *For every* $i, j \in [n], \tau \in [\lceil\log(q)\rceil]$, *define* $\{\mathsf{M}_b^{(i,j,\tau)}\}$ *as follows:*
     - *If* $F = \mathcal{Z}$, *then for every* $i \in [n], j \in [m], \tau \in [\lceil\log(q)\rceil]$, *let*
     $$\mathsf{M}_b^{(i,j,\tau)} = \mathbf{0} \in \mathbb{Z}_q^{n\times n},$$
     - *If* $F = H_r$, *then for every* $i \in [n], j \in [m], \tau \in [\lceil\log(q)\rceil]$, *let*
     $$\mathsf{M}_b^{(i,j,\tau)} = (b \oplus r_{\phi(i,j,\tau)}) \cdot \mathbf{I}_{n\times n}.$$
  2. *For every* $i \in [n], j \in [m], \tau \in [\lceil\log(q)\rceil], b \in \{0,1\}$, *compute:*
  $$pk_b^{(i,j,\tau)} = \mathbf{S}_b^{(i,j,\tau)}\mathbf{A} + \mathbf{E}_b^{(i,j,\tau)} + \mathsf{M}_b^{(i,j,\tau)},$$
  $$sk_b^{(i,j,\tau)} = \left(\left\{\mathbf{S}_b^{(i,j,\tau)}, \mathbf{E}_b^{(i,j,\tau)}\right\}\right),$$
  where for every $i \in [n], j \in [m], \tau \in [\lceil\log(q)\rceil], b \in \{0,1\}$*:*
     - $\mathbf{S}_b^{(i,j,\tau)} \leftarrow D_{\mathbb{Z}_q,\sigma}^{n\times n}$,
     - $\mathbf{E}_b^{(i,j,\tau)} \leftarrow D_{\mathbb{Z}_q,\sigma}^{n\times m}$.

3. *Output* $pk_F = \left(\mathbf{A}, \left\{pk_b^{(i,j,\tau)}\right\}_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil],b\in\{0,1\}}}\right)$ *and* $sk_F = \left\{sk_b^{(i,j,\tau)}\right\}_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil],b\in\{0,1\}}}$.

- $\mathcal{E}(pk_F, x)$*: on input* pk_F *and* $x \in \{0,1\}^\ell$*, proceed as follows:*

  1. *Parse* $pk_F = \left(\mathbf{A}\left\{pk_b^{(i,j,\tau)}\right\}_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil],b\in\{0,1\}}}\right)$
  2. *Output* $\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} pk_{x_{\phi(i,j,\tau)}}^{(i,j,\tau)}$.

Claim 16 (Correctness)**.** *Let* $(\mathcal{KG}, \mathcal{E})$ *be the pair of* PPT *algorithms in Construction 10. Let* $(pk_F, sk_F) \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$ *with* $F \in \{\mathcal{Z}\} \cup \{H_r : r \in \{0,1\}^\ell\}$*. Then, the output of* $\mathcal{E}(pk_F, x)$ *is of the form:*

$$\mathcal{E}(pk_F, x) = \mathbf{S}_x\mathbf{A} + \mathbf{E}_x + F(x),$$

where $\mathbf{S}_x \in \mathbb{Z}_q^{n\times n}$ *and* $\mathbf{E}_x \in \mathbb{Z}_q^{n\times m}$ *with* $||\mathbf{E}_x|| \le nm^2\sigma\lceil\log(q)\rceil$*. Moreover, there is an efficient algorithm that recovers* $\mathbf{S}_x$ *given* (pk_F, sk_F)*.*

Proof. Let $(pk_F, sk_F) \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$. Parse $pk_F = \left(\mathbf{A}, \left\{pk_b^{(i,j,\tau)}\right\}_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil],b\in\{0,1\}}}\right)$ and $sk_F = \left\{sk_b^{(i,j,\tau)}\right\}_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil],b\in\{0,1\}}}$, where:

$$pk_b^{(i,j,\tau)} = \mathbf{S}_b^{(i,j,\tau)}\mathbf{A} + \mathbf{E}_b^{(i,j,\tau)} + \mathsf{M}_b^{(i,j,\tau)},$$

$$sk_b^{(i,j,\tau)} = \left(\{\mathbf{S}_b^{(i,j,\tau)}, \mathbf{E}_b^{(i,j,\tau)}\}\right).$$

There are two cases to consider here:

Case 1. $F = \mathcal{Z}$: in this case, $\mathsf{M}_b^{(i,j,\tau)} = \mathbf{0}$, for every $i \in [n], j \in [m], \tau \in [\lceil\log(q)\rceil], b \in \{0,1\}$. Thus, the following holds:

$$\begin{aligned}\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} pk_{x_{\phi(i,j,\tau)}}^{(i,j,\tau)} &= \underbrace{\left(\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} \mathbf{S}_{x_{\phi(i,j,\tau)}}^{(i,j,\tau)}\right)}_{\mathbf{S}_x}\mathbf{A} + \underbrace{\left(\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} \mathbf{E}_{x_{\phi(i,j,\tau)}}^{(i,j,\tau)}\right)}_{\mathbf{E}_x} + \left(\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} \mathsf{M}_{x_{\phi(i,j,\tau)}}^{(i,j,\tau)}\right)\\ &= \mathbf{S}_x\mathbf{A} + \mathbf{E}_x + \mathcal{Z}(x).\end{aligned}$$

Moreover, we have that $||\mathbf{E}_b^{(i,j,\tau)}|| \le m\sigma$ and thus, $||\mathbf{E}_x|| \le nm^2\sigma\lceil\log(q)\rceil$.

Case 2. $F = H_r$:

$$\begin{aligned}\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} pk^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}} &= \mathbf{S}_x\mathbf{A} + \mathbf{E}_x + \left(\sum_{\substack{i\in[n],j\in[m],\\ \tau\in[\lceil\log(q)\rceil]}} \mathsf{M}^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}}\right)\\ &= \mathbf{S}_x\mathbf{A} + \mathbf{E}_x + H_r(x),\end{aligned}$$

where $\mathbf{S}_x$ and $\mathbf{E}_x$ are as defined above. The second equality holds because of the fact that $\mathsf{M}^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}}$ has the value $(b \oplus r_{\phi(i,j,\tau)}) \cdot 2^\tau$ in the $(i, j)^{th}$ position and zero, everywhere else. Thus, summing up all the $\mathsf{M}^{(i,j,\tau)}_{x_{\phi(i,j,\tau)}}$ matrices results in the matrix M, where $x \oplus r$ is the binary decomposition of M.

Finally, it is clear that $\mathbf{S}_x$ can be efficiently recovered from sk_F and x. □

Claim 17 (Shift-hiding property)**.** *Assuming the quantum hardness of learning with errors, the pair $(\mathcal{KG}, \mathcal{E})$ in Construction 10 has the property that*

$$\{pk_{\mathcal{Z}}\} \approx_c \{pk_{H_r}\},$$

for any pk_F with $(pk_F, sk_F) \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, F)$, where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, $r \in \{0,1\}^\ell$ and for any $F \in \{\mathcal{Z}, H_r\}$.

Proof. For every $i \in [n], j \in [m], \tau \in [\lceil\log(q)\rceil]$, $b \in \{0,1\}$, let $\mathsf{M}^{(i,j,\tau)}_b = (b \oplus r_{\phi(i,j,\tau)}) \cdot \mathbf{I}_{n\times n}$. Then from the quantum hardness of learning with errors, the following holds for every (i, j, τ) and $b \in \{0,1\}$:

$$\{\mathbf{S}^{(i,j,\tau)}_b\mathbf{A} + \mathbf{E}^{(i,j,\tau)}_b\} \approx_c \{\mathbf{S}^{(i,j,\tau)}_b\mathbf{A} + \mathbf{E}^{(i,j,\tau)}_b + \mathsf{M}^{(i,j,\tau)}_b\}.$$

Since $\{\mathbf{S}^{(i,j,\tau)}_b\}$ and $\{\mathbf{E}^{(i,j,\tau)}_b\}$ are sampled independently for every (i, j, τ) and $b \in \{0,1\}$, the proof of the claim follows. □

Remark 37. *When consider the all-zeroes function $\mathcal{Z}$, we drop the notation from the parameters. For instance, we denote $pk_{\mathcal{Z}}$ to be simply pk.*

Construction. We consider the following parameters which are relevant to our PRF construction. Let $n, m \in \mathbb{N}$ and let $q \in \mathbb{N}$ be a modulus with $q = 2^{o(n)}$, and let $\ell = nm\lceil\log q\rceil$. Let σ be a parameter with $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and let $p \ll q$ be a sufficiently large rounding parameter with

$$n \cdot m^3\sigma^2\lceil\log q\rceil = (q/p) \cdot 2^{-o(n)}.$$

We describe our construction below.

Construction 11 (Revocable PRF scheme)**.** *Let $n \in \mathbb{N}$ be the security parameter and $m \in \mathbb{N}$. Let $q \geq 2$ be a prime and let $\sigma > 0$ be a parameter. Let $(\mathcal{KG}, \mathcal{E})$ be the procedure in Construction 10. Our revocable PRF scheme is defined as follows:*

- $\mathsf{Gen}(1^\lambda)$*: This is the following key generation procedure:*

  1. *Sample $(\mathbf{A}, \mathsf{td}_\mathbf{A}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$.*
  2. *Compute $\kappa_\mathcal{Z} \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, \mathcal{Z})$, where $\mathcal{Z} : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$ is the such that $\mathcal{Z}(x)$ outputs an all zero matrix for every $x \in \{0,1\}^\ell$. Parse $\kappa_\mathcal{Z}$ as (pk, sk).*
  3. *Generate a Gaussian superposition $(|\psi_\mathbf{y}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ with*

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

  Output $k = (pk, sk, \mathbf{y})$, $\varrho_k = (pk, |\psi_\mathbf{y}\rangle)$ and $\mathsf{msk} = \mathsf{td}_\mathbf{A}$.

- $\mathsf{PRF}(k, x)$*: this is the following procedure:*

  1. *Parse the key k as a tuple $(pk, sk), \mathbf{y})$.*
  2. *Output $\lfloor \mathbf{S}_x \mathbf{y} \rceil_p$. Here, $\mathbf{S}_x \in \mathbb{Z}_q^{n \times n}$ is a matrix that can be efficiently recovered from sk as stated in Claim 16.*

- $\mathsf{Eval}(\varrho_k, x)$*: this is the following evaluation algorithm:*

  1. *Parse ϱ_k as (pk, ϱ).*
  2. *Compute $\mathsf{M}_x \leftarrow \mathcal{E}(pk, x)$.*
  3. *Measure the register Aux of the state $U(\varrho \otimes |0\rangle\langle 0|_\mathsf{Aux})U^\dagger$. Denote the resulting outcome to be $\mathbf{z}$, where U is defined as follows:*

  $$U |\mathbf{t}\rangle |0\rangle_\mathsf{Aux} \to |\mathbf{t}\rangle |\lfloor \mathsf{M}_x \cdot \mathbf{t} \rceil_p\rangle_\mathsf{Aux} .$$

  4. *Output $\mathbf{z}$.*

- $\mathsf{Revoke}(\mathsf{msk}, \varrho)$*: given as input the trapdoor $\mathsf{td}_\mathbf{A} \leftarrow \mathsf{msk}$, apply the projective measurement $\{|\psi_\mathbf{y}\rangle\langle\psi_\mathbf{y}|, I - |\psi_\mathbf{y}\rangle\langle\psi_\mathbf{y}|\}$ onto the state ϱ using the procedure $\mathsf{QSampGauss}(\mathbf{A}, \mathsf{td}_\mathbf{A}, \mathbf{y}, \sigma)$ in Algorithm 4. Output Valid if the measurement is successful, and $\mathsf{Invalid}$ otherwise.*

Lemma 30. *The above scheme satisfies correctness for our choice of parameters.*

Proof. The correctness of revocation follows immediately from the correctness of $\mathsf{QSampGauss}$ in Algorithm 4, which we showed in Theorem 16. Next, we show the correctness of evaluation. Let $\kappa_{\mathcal{Z}} \leftarrow \mathcal{KG}(1^n, 1^m, q, \mathbf{A}, \mathcal{Z})$ with $\kappa_{\mathcal{Z}} = (\mathsf{pk}, \mathsf{SK})$. From Claim 16, we have for any $x \in \{0,1\}^\ell$:

$$\mathcal{E}(\mathsf{pk}, x) = \mathbf{S}_x\mathbf{A} + \mathbf{E}_x \pmod q,$$

where $\mathbf{S}_x \in \mathbb{Z}_q^{n\times n}$ and $\mathbf{E}_x \in \mathbb{Z}_q^{n\times m}$ with $\|\mathbf{E}_x\| \leq nm^2\sigma\lceil\log(q)\rceil$. Recall that $\mathsf{GenGauss}(\mathbf{A},\sigma)$ outputs a state $|\psi_{\mathbf{y}}\rangle$ that is overwhelmingly supported on vectors $\mathbf{t} \in \mathbb{Z}_q^m$ such that $\|\mathbf{t}\| \leq \sigma\sqrt{\frac{m}{2}}$ with $\mathbf{A}\cdot\mathbf{t} = \mathbf{y} \pmod q$. Therefore, we have for any input $x \in \{0,1\}^\ell$:

$$\lfloor\mathcal{E}(\mathsf{pk}, x)\cdot\mathbf{t}\rceil_p = \lfloor\mathbf{S}_x\mathbf{A}\cdot\mathbf{t} + \mathbf{E}_x\cdot\mathbf{t}\rceil_p = \lfloor\mathbf{S}_x\cdot\mathbf{y} + \mathbf{E}_x\cdot\mathbf{t}\rceil_p = \lfloor\mathbf{S}_x\cdot\mathbf{y}\rceil_p\,,$$

where the last equality follows from the fact that

$$\|\mathbf{E}_x\cdot\mathbf{t}\|_2 \leq \|\mathbf{E}_x\|_2\cdot\|\mathbf{t}\|_2 \leq \sqrt{m}\cdot\|\mathbf{E}_x\|\cdot\|\mathbf{t}\|_2 \leq n\sqrt{m}m^2\sigma\lceil\log(q)\rceil\cdot\sigma\sqrt{m/2}.$$

and $n\cdot m^3\sigma^2\lceil\log q\rceil = (q/p)\cdot 2^{-o(n)}$ for our choice of parameters. □

Proof of security. Our first result on the security of Construction 11 concerns $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda), 1)$-security, i.e., we assume that revocation succeeds with overwhelming probability.

Theorem 38. *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and $m \geq 2n\log q$, each parameterized by $\lambda \in \mathbb{N}$. Let $\ell = nm\lceil\log q\rceil$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and $\alpha \in (0,1)$ be any noise ratio with $1/\alpha = \sigma\cdot 2^{o(n)}$, and let $p \ll q$ be a sufficiently large rounding parameter with*

$$n\cdot m^3\sigma^2\lceil\log q\rceil = (q/p)\cdot 2^{-o(n)}.$$

Then, assuming the quantum subexponential hardness of $\mathsf{LWE}^m_{n,q,\alpha q}$ and $\mathsf{SIS}^m_{n,q,\sigma\sqrt{2m}}$, our revocable PRF scheme $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ defined in Construction 11 satisfies $(\mathsf{negl}(\lambda), \mathsf{negl}(\lambda), 1)$-revocation security according to Definition 44.

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

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

for some function $\varepsilon(\lambda)$ with respect to security experiment $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, b)$ from Figure 5.14. To complete the proof, it suffices to show that $\varepsilon(\lambda)$ is negligible.

Suppose for the sake of contradition that $\epsilon(\lambda) = 1/\mathsf{poly}(\lambda)$. Let us now introduce a sequence of hybrid experiments which will be relevant for the remainder of the proof.

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

Initialization Phase:

- The challenger runs the procedure $\mathsf{Gen}(1^\lambda)$:

  1. Sample $(\mathbf{A}, \mathsf{td}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^n, 1^m, q)$.
  2. Generate $\mathbf{A}_N \in \mathbb{Z}_q^{(n+m)\times m}$ with $\overline{\mathbf{A}_N} \xleftarrow{\$} \mathbb{Z}_q^{m\times m}$ and $\underline{\mathbf{A}_N} = \mathbf{A}$.
  3. Compute $\kappa_{\mathcal{Z}} \leftarrow \mathcal{KG}(1^n, 1^m, 1^q, \mathbf{A}_N, \mathcal{Z})$, where $\mathcal{KG}$ is as defined in Construction 10 and $\mathcal{Z} : \{0,1\}^\ell \to \mathbb{Z}_q^{n\times m}$ is such that $\mathcal{Z}(x)$ outputs an all zero matrix for every $x \in \{0,1\}^\ell$. Parse $\kappa_{\mathcal{Z}}$ as (pk, sk).
  4. Generate $(|\psi_{\mathbf{y}}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ with

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

  5. Let $k = (pk, sk, \mathbf{y})$, $\varrho_k = (pk, |\psi_{\mathbf{y}}\rangle)$ and $\mathsf{msk} = \mathsf{td}_{\mathbf{A}}$.

- The challenger sends $\varrho_k = (pk, |\psi_{\mathbf{y}}\rangle)$ to $\mathcal{A}$.

Revocation Phase:

- The challenger sends the message REVOKE to $\mathcal{A}$.
- $\mathcal{A}$ generates a (possibly entangled) bipartite quantum state $\varrho_{R,\text{AUX}}$ in systems $\mathcal{H}_R \otimes \mathcal{H}_{\text{AUX}}$ with $\mathcal{H}_R = \mathcal{H}_q^m$, returns system R and holds onto the auxiliary system AUX.
- The challenger runs $\mathsf{Revoke}(\mathsf{msk}, \varrho_R)$, where ϱ_R is the reduced state in system R. If the outcome is $\mathsf{Invalid}$, the challenger aborts.

Guessing Phase:

- The challenger samples $x \leftarrow \{0,1\}^\ell$ and sends (x, y) to $\mathcal{A}$, where
  - If $b = 0$: compute $\mathbf{S}_x$ from sk as in Claim 16. Set $y = \lfloor \mathbf{S}_x \mathbf{y} \rceil_p$.
  - If $b = 1$: sample $y \leftarrow \{0,1\}^n$.
- $\mathcal{A}$ outputs a string b' and wins if $b' = b$.

Figure 5.14: The revocable PRF experiment $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, b)$ for Construction 11.

Let $\mathsf{RevDual} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Revoke})$ be the n-bit key-revocable Dual-Regev scheme from Construction 8. Fix $\mu = 0^n$, where μ is the challenge message in the dual-Regev encryption security.

H_0: This is $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, 0)$ in Figure 5.14.

H_1: This is the same experiment as $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, 0)$, except for the following changes:

- Sample a random string $r \leftarrow \{0,1\}^\ell$.
- Run the procedure $\mathsf{RevDual.KeyGen}(1^\lambda)$ instead of $\mathsf{GenTrap}(1^n, 1^m, q)$ and $\mathsf{GenGauss}(\mathbf{A}, \sigma)$ to obtain $(\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathbf{y} \in \mathbb{Z}_q^n, \mathsf{msk}, \mathsf{sk})$.
- Compute $(\mathsf{CT}_1, \mathsf{CT}_2) \leftarrow \mathsf{RevDual.Enc}(\mathbf{A}, \mathbf{y}, \mu)$, where $\mathsf{CT}_1 \in \mathbb{Z}_q^{n\times m}$ and $\mathsf{CT}_2 \in \mathbb{Z}_q^n$.
- Set $x = r \oplus \mathsf{bindecomp}(\mathsf{CT}_1)$.

The rest of the hybrid is the same as before.

Note that Hybrids H_0 and H_1 are identically distributed.

H_2: This is the same experiment as before, except that the challenger now uses an alternative key-generation algorithm:

- As before, run the procedure $\mathsf{RevDual.KeyGen}(1^\lambda)$ instead of $\mathsf{GenTrap}(1^n, 1^m, q)$ and $\mathsf{GenGauss}(\mathbf{A}, \sigma)$ to obtain $(\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathbf{y} \in \mathbb{Z}_q^n, \mathsf{msk}, \mathsf{sk})$. Sample $r \leftarrow \{0,1\}^\ell$.
- Let $H_r : \{0,1\}^\ell \to \mathbb{Z}_q^{n\times m}$ be as defined in the beginning of Section 5.7.
- Run the alternate algorithm $\kappa_H \leftarrow \mathcal{KG}(1^n, 1^m, 1^q, \mathbf{A}, H_r)$ instead of $\kappa_{\mathcal{Z}} \leftarrow \mathcal{KG}(1^n, 1^m, 1^q, \mathbf{A}, \mathcal{Z})$.
- Compute the ciphertext $(\mathsf{CT}_1^*, \mathsf{CT}_2^*) \leftarrow \mathsf{RevDual.Enc}(\mathbf{A}, \mathbf{y}, \mu)$, where $\mathsf{CT}_1^* \in \mathbb{Z}_q^{n\times m}$. Then, set $x^* = r \oplus \mathsf{bindecomp}(\mathsf{CT}_1^*)$. Send x^* to the adversary in the guessing phase.

H_3: This is the same hybrid as before, except that we choose $\mathsf{CT}_1^* \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and $\mathsf{CT}_2^* \xleftarrow{\$} \mathbb{Z}_q^n$.

H_4: This is the $\mathsf{Expt}^{\mathcal{A}}(1^\lambda, 1)$ in Figure 5.14.

Note that hybrids H_3 and H_4 are identically distributed.

We now show the following.

Claim 18. *By the shift-hiding property of $(\mathcal{KG}, \mathcal{E})$ in Claim 17, we have that the two hybrids* H_1 *and* H_2 *are computationally indistinguishable, i.e.,*

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

Proof. Suppose for the sake of contradiction that there exists a non-negligble difference in the advantage of the adversary $\mathcal{A}$ in the two hybrids H_1 and H_2.

We now design a reduction $\mathcal{B}$ that violates the shift-hiding property as follows.

1. Sample $r \xleftarrow{\$} \{0,1\}^\ell$. Send $(\mathcal{Z}, H_r)$ to the challenger.
2. The challenger responds with $pk = \left(\mathbf{A}, \left\{\mathsf{CT}_b^{(i,j,\tau)}\right\}_{\substack{i\in[n],j\in[m] \\ \tau\in[\lceil\log(q)\rceil],b\in\{0,1\}}}\right)$
3. Compute $(|\psi_\mathbf{y}\rangle, \mathbf{y} \in \mathbb{Z}_q^n) \leftarrow \mathsf{GenGauss}(\mathbf{A}, \sigma)$ from the challenger.
4. Set $\varrho_k = (pk, \varrho)$.
5. Compute $(\mathsf{CT}_1, \mathsf{CT}_2) \leftarrow \mathsf{RevDual.Enc}(\mathbf{A}, \mathbf{y}, \mu)$, where $\mathsf{CT}_1 \in \mathbb{Z}_q^{n\times m}$ and $\mathsf{CT}_2 \in \mathbb{Z}_q^n$. Then, set $x^* = r \oplus \mathsf{bindecomp}(\mathsf{CT}_1)$.
6. Compute $\mathsf{Eval}(\varrho_k, x^*)$ to obtain y^* while recovering ϱ_k^* (using the "Almost as Good As New" Lemma, Lemma 1) such that $\mathsf{T}D(\varrho_k^*, \varrho_k) \leq \mathsf{negl}(\lambda)$.
7. Send ϱ_k^* to $\mathcal{A}$.
8. $\mathcal{A}$ computes a state on two registers R and AUX. It returns the state on the register R.
9. $\mathcal{A}$, on input the register AUX and (x^*, y^*), outputs a bit b'.
10. Output b'.

If pk is generated using $\mathcal{KG}(1^n, 1^m, q, \mathbf{A}, \mathcal{Z})$ then we are precisely in H_1 (except that Revoke is not performed). Moreover, if pk is generated using $\mathcal{KG}(1^n, 1^m, q, \mathbf{A}, H_r)$, then we are in hybrid H_2 (except that Revoke is not performed). Therefore, the reduction $\mathcal{B}$ has a non-negligible advantage at distinguishing H_1 and H_2 whenever Revoke outputs $\top$ on system R. By Lemma 28, this implies that we can use $\mathcal{B}$ to break the shift-hiding property with non-negligible advantage – thereby yielding a contradiction. This proves the claim. □

Next, we invoke the security of the n-bit variant of our key-revocable Dual-Regev scheme (which is implied by Theorem 24) to show the following.

Claim 19. *By the security of our n-bit key-revocable Dual-Regev encryption scheme, we have that the two hybrids* H_2 *and* H_3 *are computationally indistinguishable, i.e.,*

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

Proof. Suppose for the sake of contradiction that there exists a non-negligble difference in the advantage of $\mathcal{A}$ in H_2 and H_3. Using $\mathcal{A}$, we can now design a reduction $\mathcal{B}$ that violates the revocation security of our n-bit revocable Dual-Regev scheme which is implicit in Theorem 25.

The reduction $\mathcal{B}$ proceeds as follows.

1. First, it receives as input $\mathbf{A}, \mathbf{y}$ and a quantum state

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

2. The reduction generates a quantum state ϱ_k as follows:

   - Sample a random string $r \overset{\$}{\leftarrow} \{0,1\}^\ell$.
   - Let $H_r : \{0,1\}^\ell \to \mathbb{Z}_q^{n \times m}$ be as defined in the beginning of Section 5.7.
   - Run the algorithm $\kappa_H \leftarrow \mathcal{KG}(1^n, 1^m, 1^q, \mathbf{A}, H_r)$ and parse κ_H as (pk, sk).
   - Set $\varrho_k = (pk, |\psi_{\mathbf{y}}\rangle)$.

   Send ϱ_k to $\mathcal{A}$.

3. $\mathcal{A}$ outputs a state on two registers R and $\textsc{aux}$. The register R is returned. The reduction forwards the register R to the challenger.

4. The reduction then gets the challenge ciphertext $\mathsf{CT} = [\mathsf{CT}_1, \mathsf{CT}_2]^\intercal \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$. The reduction then sets

$$x^* = r \oplus \mathsf{bindecomp}(\mathsf{CT}_1)$$

   and sends x^* to $\mathcal{A}$ in the guessing phase, together with $y = \lfloor \mathbf{S}_{x^*}\mathbf{y} + \mathsf{CT}_2 \rceil_p$ which is computed using the secret key SK (c.f. Claim 16).

5. $\mathcal{A}$ outputs a bit b'. $\mathcal{B}$ outputs b'.

There are two cases to consider here. In the first case, we have $\mathsf{CT} = [\mathsf{CT}_1, \mathsf{CT}_2]^\intercal \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$ is a Dual-Regev ciphertext. Here, $y = \lfloor \mathbf{S}_{x^*}\mathbf{y} + \mathsf{CT}_2 \rceil_p$ precisely corresponds to the output of the pseudorandom function on ϱ_k and x. In the second case, we have $\mathsf{CT} = [\mathsf{CT}_1, \mathsf{CT}_2]^\intercal$, where

$\mathsf{CT}_1 \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ and $\mathsf{CT}_2 \xleftarrow{\$} \mathbb{Z}_q^{m}$. Therefore, the resulting string $y = \lfloor \mathbf{S}_{x^*}\mathbf{y} + \mathsf{CT}_2 \rceil_p$ is negligibly close (in total variation distance) to a uniform distribution on $\mathbb{Z}_p^m$.

Putting everything together, we find that the first case corresponds precisely to H_2, whereas the second case corresponds to H_3. As a result, $\mathcal{B}$ violates the revocation security of our n-bit revocable Dual-Regev scheme which is implicit in Theorem 24.This completes the proof. □

Putting everything together, we have shown that

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

□

Finally, we prove that Construction 11 achieves the stronger notion of $(\mathsf{negl}(\lambda), 1 - 1/\mathsf{poly}(\lambda), 1)$-security assuming Conjecture 1.

Theorem 39. *Let $n \in \mathbb{N}$ and q be a prime modulus with $q = 2^{o(n)}$ and $m \geq 2n \log q$, each parameterized by $\lambda \in \mathbb{N}$. Let $\ell = nm\lceil \log q \rceil$. Let $\sqrt{8m} < \sigma < q/\sqrt{8m}$ and $\alpha \in (0,1)$ be any noise ratio with $1/\alpha = \sigma \cdot 2^{o(n)}$, and let $p \ll q$ be a sufficiently large rounding parameter with*

$$n \cdot m^3 \sigma^2 \lceil \log q \rceil = (q/p) \cdot 2^{-o(n)}.$$

Then, assuming Conjecture 1, our revocable PRF *scheme* $(\mathsf{Gen}, \mathsf{PRF}, \mathsf{Eval}, \mathsf{Revoke})$ *defined in Construction 11 has* $(\mathsf{negl}(\lambda), 1 - 1/\mathsf{poly}(\lambda), 1)$*-revocation security according to Definition 44.*

Proof. The proof is the same as in Theorem 38, except that we invoke Theorem 25 instead of Theorem 24 for Dual-Regev security. □

BIBLIOGRAPHY

[1] Scott Aaronson. *How Much Structure Is Needed for Huge Quantum Speedups?* 2022. DOI: 10.48550/ARXIV.2209.06930. URL: https://arxiv.org/abs/2209.06930.

[2] Scott Aaronson. “Quantum copy-protection and quantum money”. In: *2009 24th Annual IEEE Conference on Computational Complexity*. IEEE. 2009, pp. 229–242.

[3] Scott Aaronson. *The Complexity of Quantum States and Transformations: From Quantum Money to Black Holes*. 2016. arXiv: 1607.05256 [quant-ph].

[4] Scott Aaronson and Paul Christiano. *Quantum Money from Hidden Subspaces*. 2012. DOI: 10.48550/ARXIV.1203.4740. URL: https://arxiv.org/abs/1203.4740.

[5] Scott Aaronson, Jiahui Liu, and Ruizhe Zhang. “Quantum Copy-Protection from Hidden Subspaces”. In: *arXiv preprint arXiv:2004.09674* (2020).

[6] Scott Aaronson et al. “New approaches for quantum copy-protection”. In: *Annual International Cryptology Conference*. Springer. 2021, pp. 526–555.

[7] Mark Adcock and Richard Cleve. “A quantum Goldreich-Levin theorem with cryptographic applications”. In: *STACS 2002*. Ed. by Helmut Alt and Afonso Ferreira. Springer, 2002, pp. 323–334. ISBN: 978-3-540-45841-8. DOI: 10.1007/3-540-45841-7_26. arXiv: quant-ph/0108095.

[8] Shweta Agrawal et al. “Public Key Encryption with Secure Key Leasing”. In: *arXiv preprint arXiv:2302.11663* (2023).

[9] Miklós Ajtai. “Generating Hard Instances of Lattice Problems (Extended Abstract)”. In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*. Ed. by Gary L. Miller. ACM, 1996, pp. 99–108. DOI: 10.1145/237814.237838. URL: https://doi.org/10.1145/237814.237838.

[10] Gorjan Alagic et al. “Computational Security of Quantum Encryption”. In: *Information Theoretic Security* (2016), pp. 47–71. ISSN: 1611-3349. DOI: 10.1007/978-3-319-49175-2_3. URL: http://dx.doi.org/10.1007/978-3-319-49175-2_3.

[11] Gorjan Alagic et al. “On Quantum Chosen-Ciphertext Attacks and Learning with Errors”. In: *Cryptography* 4.1 (2020). ISSN: 2410-387X. DOI: 10.3390/cryptography4010010. URL: https://www.mdpi.com/2410-387X/4/1/10.

[12] Andris Ambainis, Mike Hamburg, and Dominique Unruh. “Quantum security proofs using semi-classical oracles”. In: *Annual International Cryptology Conference*. Springer. 2019, pp. 269–295.

[13] Prabhanjan Ananth and Fatih Kaleoglu. “A Note on Copy-Protection from Random Oracles”. In: *arXiv preprint arXiv:2208.12884* (2022).

[14] Prabhanjan Ananth and Fatih Kaleoglu. *Unclonable Encryption, Revisited*. 2021. DOI: 10.48550/ARXIV.2103.15009. URL: https://arxiv.org/abs/2103.15009.

[15] Prabhanjan Ananth and Fatih Kaleoglu. "Unclonable encryption, revisited". In: *Theory of Cryptography: 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8–11, 2021, Proceedings, Part I*. Springer. 2021, pp. 299–329.

[16] Prabhanjan Ananth, Fatih Kaleoglu, and Qipeng Liu. "Cloning Games: A General Framework for Unclonable Primitives". In: *arXiv preprint arXiv:2302.01874* (2023).

[17] Prabhanjan Ananth and Rolando L La Placa. "Secure Software Leasing". In: *Eurocrypt* (2021).

[18] Prabhanjan Ananth et al. *On the Feasibility of Unclonable Encryption, and More*. Cryptology ePrint Archive, Paper 2022/884. https://eprint.iacr.org/2022/884. 2022. URL: https://eprint.iacr.org/2022/884.

[19] Saikrishna Badrinarayanan et al. "Post-zeroizing Obfuscation: New Mathematical Tools, and the Case of Evasive Circuits". In: *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*. 2016, pp. 764–791.

[20] W. Banaszczyk. "New bounds in some transference theorems in the geometry of numbers." In: *Mathematische Annalen* 296.4 (1993), pp. 625–636. URL: http://eudml.org/doc/165105.

[21] Abhishek Banerjee, Chris Peikert, and Alon Rosen. "Pseudorandom functions and lattices". In: *Advances in Cryptology–EUROCRYPT 2012: 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings 31*. Springer. 2012, pp. 719–737.

[22] Boaz Barak et al. "On the (Im)Possibility of Obfuscating Programs". In: *J. ACM* 59.2 (May 2012). ISSN: 0004-5411. DOI: 10.1145/2160158.2160159. URL: https://doi.org/10.1145/2160158.2160159.

[23] James Bartusek and Dakshita Khurana. *Cryptography with Certified Deletion*. 2022. DOI: 10.48550/ARXIV.2207.01754. URL: https://arxiv.org/abs/2207.01754.

[24] James Bartusek, Dakshita Khurana, and Alexander Poremba. *Publicly-Verifiable Deletion via Target-Collapsing Functions*. 2023. arXiv: 2303.08676 [quant-ph].

[25] James Bartusek et al. *Obfuscation and Outsourced Computation with Certified Deletion*. Cryptology ePrint Archive, Paper 2023/265. https://eprint.iacr.org/2023/265. 2023. URL: https://eprint.iacr.org/2023/265.

[26] Mihir Bellare and Phillip Rogaway. "Random Oracles Are Practical: A Paradigm for Designing Efficient Protocols". In: *Proceedings of the 1st ACM Conference on Computer and Communications Security*. CCS '93. Fairfax, Virginia, USA: Association for Computing Machinery, 1993, pp. 62–73. ISBN: 0897916298. DOI: 10.1145/168588.168596. URL: https://doi.org/10.1145/168588.168596.

[27] Shalev Ben-David and Or Sattath. "Quantum Tokens for Digital Signatures". In: (2016). DOI: 10.48550/ARXIV.1609.09047. URL: https://arxiv.org/abs/1609.09047.

[28] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. “Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation”. In: *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*. STOC ’88. Chicago, Illinois, USA: Association for Computing Machinery, 1988, pp. 1–10. ISBN: 0897912640. DOI: 10.1145/62212.62213. URL: https://doi.org/10.1145/62212.62213.

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

[30] Charles H. Bennett et al. “Strengths and Weaknesses of Quantum Computing”. In: *SIAM J. Comput.* 26.5 (Oct. 1997), pp. 1510–1523. ISSN: 0097-5397. DOI: 10.1137/S0097539796300933. URL: https://doi.org/10.1137/S0097539796300933.

[31] Nina Bindel et al. “Tighter Proofs of CCA Security in the Quantum Random Oracle Model”. In: *Theory of Cryptography*. Ed. by Dennis Hofheinz and Alon Rosen. Cham: Springer International Publishing, 2019, pp. 61–90. ISBN: 978-3-030-36033-7.

[32] Nir Bitansky, Zvika Brakerski, and Yael Tauman Kalai. *Constructive Post-Quantum Reductions*. 2022. DOI: 10.48550/ARXIV.2203.02314. URL: https://arxiv.org/abs/2203.02314.

[33] Dan Boneh et al. *Fully Key-Homomorphic Encryption, Arithmetic Circuit ABE, and Compact Garbled Circuits*. Cryptology ePrint Archive, Paper 2014/356. https://eprint.iacr.org/2014/356. 2014. URL: https://eprint.iacr.org/2014/356.

[34] Dan Boneh et al. “Random oracles in a quantum world”. In: *International Conference on the Theory and Application of Cryptology and Information Security*. Springer. 2011, pp. 41–69.

[35] Raphael Bost et al. “Machine Learning Classification over Encrypted Data”. In: *IACR Cryptology ePrint Archive* 2014 (2014), p. 331. URL: https://eprint.iacr.org/2014/331.

[36] Zvika Brakerski. *Quantum FHE (Almost) As Secure As Classical*. Cryptology ePrint Archive, Report 2018/338. https://ia.cr/2018/338. 2018.

[37] 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. USA: IEEE Computer Society, 2011, pp. 97–106. ISBN: 9780769545714. DOI: 10.1109/FOCS.2011.12. URL: https://doi.org/10.1109/FOCS.2011.12.

[38] Zvika Brakerski et al. *A Cryptographic Test of Quantumness and Certifiable Randomness from a Single Quantum Device*. 2021. arXiv: 1804.00640 [quant-ph].

[39] Zvika Brakerski et al. “Factoring and pairings are not necessary for io: Circular-secure lwe suffices”. In: *Cryptology ePrint Archive* (2020).

[40] Anne Broadbent and Rabib Islam. “Quantum encryption with certified deletion”. In: *Theory of Cryptography Conference*. Springer. 2020, pp. 92–122.

[41] Anne Broadbent and Sébastien Lord. “Uncloneable Quantum Encryption via Oracles”. In: *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*. Ed. by Steven T. Flammia. Vol. 158. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, 4:1–4:22. DOI: 10.4230/LIPIcs.TQC.2020.4.

[42] Anne Broadbent et al. “Secure Software Leasing Without Assumptions”. In: *Theory of Cryptography*. Springer International Publishing, 2021, pp. 90–120. DOI: 10.1007/978-3-030-90459-3_4. URL: https://doi.org/10.1007%2F978-3-030-90459-3_4.

[43] Anne Broadbent et al. “Secure software leasing without assumptions”. In: *Theory of Cryptography Conference*. Springer. 2021, pp. 90–120.

[44] Ran Canetti, Oded Goldreich, and Shai Halevi. “The Random Oracle Methodology, Revisited”. In: *J. ACM* 51.4 (July 2004), pp. 557–594. ISSN: 0004-5411. DOI: 10.1145/1008731.1008734. URL: https://doi.org/10.1145/1008731.1008734.

[45] Ran Canetti et al. “Adaptively Secure Multi-Party Computation”. In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC ’96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, pp. 639–648. ISBN: 0897917855. DOI: 10.1145/237814.238015. URL: https://doi.org/10.1145/237814.238015.

[46] Ran Canetti et al. “On Symmetric Encryption and Point Obfuscation”. In: *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*. Vol. 5978. Lecture Notes in Computer Science. Springer, 2010, pp. 52–71. DOI: 10.1007/978-3-642-11799-2_4. URL: https://www.iacr.org/archive/tcc2010/59780052/59780052.pdf.

[47] Yudong Cao et al. “Quantum Chemistry in the Age of Quantum Computing”. In: *Chemical Reviews* 119.19 (2019). PMID: 31469277, pp. 10856–10915. DOI: 10.1021/acs.chemrev.8b00803. eprint: https://doi.org/10.1021/acs.chemrev.8b00803. URL: https://doi.org/10.1021/acs.chemrev.8b00803.

[48] Yilei Chen, Qipeng Liu, and Mark Zhandry. *Quantum Algorithms for Variants of Average-Case Lattice Problems via Filtering*. 2021. arXiv: 2108.11015 [quant-ph].

[49] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. “GGH15 beyond permutation branching programs: proofs, attacks, and candidates”. In: *Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part II 38*. Springer. 2018, pp. 577–607.

[50] Arka Rai Choudhuri, Abhihsek Jain, and Zhengzhong Jin. “Snargs for P from LWE”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 68–79.

[51] Xavier Coiteux-Roy and Stefan Wolf. “Proving Erasure”. In: *2019 IEEE International Symposium on Information Theory (ISIT)* (July 2019). DOI: 10.1109/isit.2019.8849661. URL: http://dx.doi.org/10.1109/ISIT.2019.8849661.

[52] Andrea Coladangelo, Christian Majenz, and Alexander Poremba. *Quantum copy-protection of compute-and-compare programs in the quantum random oracle model*. 2020. arXiv: 2009.13865 [quant-ph].

[53] Andrea Coladangelo et al. *Hidden Cosets and Applications to Unclonable Cryptography*. 2021. arXiv: 2107.05692 [cs.CR].

[54] Andrea Coladangelo et al. "Hidden cosets and applications to unclonable cryptography". In: *Annual International Cryptology Conference*. Springer. 2021, pp. 556–584.

[55] Lalita Devadas et al. "Succinct LWE sampling, random polynomials, and obfuscation". In: *Theory of Cryptography: 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8–11, 2021, Proceedings, Part II 19*. Springer. 2021, pp. 256–287.

[56] DGBJ Dieks. "Communication by EPR devices". In: *Physics Letters A* 92.6 (1982), pp. 271–272.

[57] Yevgeniy Dodis et al. "Public-Key Encryption Schemes with Auxiliary Inputs". In: *Theory of Cryptography*. Ed. by Daniele Micciancio. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 361–381. ISBN: 978-3-642-11799-2.

[58] Fré déric Dupuis, Jesper Buus Nielsen, and Louis Salvail. "Secure Two-Party Quantum Evaluation of Unitaries against Specious Adversaries". In: *Advances in Cryptology – CRYPTO 2010*. Springer Berlin Heidelberg, 2010, pp. 685–706. DOI: 10.1007/978-3-642-14623-7_37. URL: https://doi.org/10.1007%2F978-3-642-14623-7_37.

[59] Stefan Dziembowski, Tomasz Kazana, and Daniel Wichs. "One-Time Computable Self-Erasing Functions". In: *Theory of Cryptography - 8th Theory of Cryptography Conference, TCC 2011*. Vol. 6597. Lecture Notes in Computer Science. Springer, 2011, p. 125. DOI: 10.1007/978-3-642-19571-6_9. URL: https://www.iacr.org/archive/tcc2011/65970124/65970124.pdf.

[60] Edward Eaton and Fang Song. "A Note on the Instantiability of the Quantum Random Oracle". In: *Post-Quantum Cryptography*. Ed. by Jintai Ding and Jean-Pierre Tillich. Cham: Springer International Publishing, 2020, pp. 503–523. ISBN: 978-3-030-44223-1.

[61] Honghao Fu and Carl A. Miller. "Local randomness: Examples and application". In: *Physical Review A* 97.3 (Mar. 2018). ISSN: 2469-9934. DOI: 10.1103/physreva.97.032324. URL: http://dx.doi.org/10.1103/PhysRevA.97.032324.

[62] Jingliang Gao. "Quantum union bounds for sequential projective measurements". In: *Physical Review A* 92.5 (2015), p. 052331.

[63] Sanjam Garg, Shafi Goldwasser, and Prashant Nalini Vasudevan. "Formalizing Data Deletion in the Context of the Right to be Forgotten". In: *IACR Cryptol. ePrint Arch.* (2020), p. 254. URL: https://eprint.iacr.org/2020/254.

[64] Sanjam Garg et al. "On the Implausibility of Differing-Inputs Obfuscation and Extractable Witness Encryption with Auxiliary Input". In: *Algorithmica* 79.4 (Dec. 2017), pp. 1353–1373. ISSN: 0178-4617. DOI: 10.1007/s00453-017-0276-6. URL: https://doi.org/10.1007/s00453-017-0276-6.

[65] Romain Gay and Rafael Pass. “Indistinguishability obfuscation from circular security”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 736–749.

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

[67] Craig Gentry. “Fully homomorphic encryption using ideal lattices”. In: *Proceedings of the forty-first annual ACM symposium on Theory of computing*. 2009, pp. 169–178.

[68] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. *Trapdoors for Hard Lattices and New Cryptographic Constructions*. Cryptology ePrint Archive, Report 2007/432. https://eprint.iacr.org/2007/432. 2007.

[69] Craig Gentry, Amit Sahai, and Brent Waters. *Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based*. Cryptology ePrint Archive, Report 2013/340. https://ia.cr/2013/340. 2013.

[70] Marios Georgiou and Mark Zhandry. “Unclonable decryption keys”. In: *Cryptology ePrint Archive* (2020).

[71] O. Goldreich and L. A. Levin. “A Hard-Core Predicate for All One-Way Functions”. In: *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*. STOC ’89. Seattle, Washington, USA: Association for Computing Machinery, 1989, pp. 25–32. ISBN: 0897913078. DOI: 10.1145/73007.73010. URL: https://doi.org/10.1145/73007.73010.

[72] Oded Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Cambridge University Press, 2006.

[73] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. “How to construct random functions”. In: *Journal of the ACM* 33.4 (1986), pp. 792–807. ISSN: 0004-5411. DOI: 10.1145/6490.6503.

[74] S Goldwasser, S Micali, and C Rackoff. “The Knowledge Complexity of Interactive Proof-Systems”. In: *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*. STOC ’85. Providence, Rhode Island, USA: Association for Computing Machinery, 1985, pp. 291–304. ISBN: 0897911512. DOI: 10.1145/22145.22178. URL: https://doi.org/10.1145/22145.22178.

[75] Daniel Gottesman. “Uncloneable encryption”. In: *Quantum Information & Computation* 3.6 (2003), pp. 581–602.

[76] Rishab Goyal, Venkata Koppula, and Brent Waters. “Lockable obfuscation”. In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2017, pp. 612–621.

[77] Alex B. Grilo, Iordanis Kerenidis, and Timo Zijlstra. “Learning-with-errors problem is easy with quantum samples”. In: *Physical Review A* 99.3 (Mar. 2019). ISSN: 2469-9934. DOI: 10.1103/physreva.99.032314. URL: http://dx.doi.org/10.1103/PhysRevA.99.032314.

[78] Lov K. Grover and Terry Rudolph. "Creating superpositions that correspond to efficiently integrable probability distributions". In: *arXiv: Quantum Physics* (2002).

[79] L. Hales and S. Hallgren. "An improved quantum Fourier transform algorithm and applications". In: *Proceedings 41st Annual Symposium on Foundations of Computer Science*. 2000, pp. 515–525. DOI: 10.1109/SFCS.2000.892139.

[80] Johan Håstad et al. "Pseudo-Random Generation from One-Way Functions". In: *PROC. 20TH STOC*. 1988, pp. 12–24.

[81] Carmit Hazay and Yehuda Lindell. *A Note on the Relation between the Definitions of Security for Semi-Honest and Malicious Adversaries*. Cryptology ePrint Archive, Paper 2010/551. https://eprint.iacr.org/2010/551. 2010. URL: https://eprint.iacr.org/2010/551.

[82] Taiga Hiroka et al. *Certified Everlasting Zero-Knowledge Proof for QMA*. 2021. arXiv: 2109.14163 [quant-ph].

[83] Taiga Hiroka et al. *Quantum Encryption with Certified Deletion, Revisited: Public Key, Attribute-Based, and Classical Communication*. 2021. arXiv: 2105.05393 [quant-ph].

[84] Taiga Hiroka et al. "Quantum Encryption with Certified Deletion, Revisited: Public Key, Attribute-Based, and Classical Communication". In: *Lecture Notes in Computer Science*. Springer International Publishing, 2021, pp. 606–636. DOI: 10.1007/978-3-030-92062-3_21. URL: https://doi.org/10.1007%2F978-3-030-92062-3_21.

[85] Intercept. "How Spies Stole The Keys To The Encryption Castle". In: *https://theintercept.com/2015/02/19/great-sim-heist/*. 2015.

[86] Aayush Jain, Huijia Lin, and Amit Sahai. "Indistinguishability obfuscation from well-founded assumptions". In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 60–73.

[87] Stanisław Jarecki and Anna Lysyanskaya. "Adaptively Secure Threshold Cryptography: Introducing Concurrency, Removing Erasures". In: *Proceedings of the 19th International Conference on Theory and Application of Cryptographic Techniques*. EUROCRYPT'00. Bruges, Belgium: Springer-Verlag, 2000, pp. 221–242. ISBN: 3540675175.

[88] Fuyuki Kitagawa and Ryo Nishimaki. "Functional Encryption with Secure Key Leasing". In: *ASIACRYPT*. 2022.

[89] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. "Secure Software Leasing from Standard Assumptions". In: *Theory of Cryptography*. Springer International Publishing, 2021, pp. 31–61. DOI: 10.1007/978-3-030-90459-3_2. URL: https://doi.org/10.1007%2F978-3-030-90459-3_2.

[90] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. "Secure software leasing from standard assumptions". In: *Theory of Cryptography Conference*. Springer. 2021, pp. 31–61.

[91] Fuyuki Kitagawa et al. *Adaptively Secure and Succinct Functional Encryption: Improving Security and Efficiency, Simultaneously*. Cryptology ePrint Archive, Paper 2018/974. https://eprint.iacr.org/2018/974. 2018. URL: https://eprint.iacr.org/2018/974.

[92] Robert Konig, Renato Renner, and Christian Schaffner. "The Operational Meaning of Min- and Max-Entropy". In: *IEEE Transactions on Information Theory* 55.9 (Sept. 2009), pp. 4337–4347. DOI: 10.1109/tit.2009.2025545. URL: https://doi.org/10.1109%2Ftit.2009.2025545.

[93] Veronika Kuchta et al. "Measure-Rewind-Measure: Tighter Quantum Random Oracle Model Proofs for One-Way to Hiding and CCA Security". In: *Advances in Cryptology – EUROCRYPT 2020*. Ed. by Anne Canteaut and Yuval Ishai. Cham: Springer International Publishing, 2020, pp. 703–728. ISBN: 978-3-030-45727-3.

[94] Seunghoon Lee et al. *Is there evidence for exponential quantum advantage in quantum chemistry?* 2022. DOI: 10.48550/ARXIV.2208.02199. URL: https://arxiv.org/abs/2208.02199.

[95] Nathalie P. de Leon et al. "Materials challenges and opportunities for quantum computing hardware". In: *Science* 372.6539 (2021), eabb2823. DOI: 10.1126/science.abb2823. eprint: https://www.science.org/doi/pdf/10.1126/science.abb2823. URL: https://www.science.org/doi/abs/10.1126/science.abb2823.

[96] Jiahui Liu et al. *Collusion Resistant Copy-Protection for Watermarkable Functionalities*. Cryptology ePrint Archive, Paper 2022/1429. https://eprint.iacr.org/2022/1429. 2022. URL: https://eprint.iacr.org/2022/1429.

[97] Qipeng Liu and Mark Zhandry. *Revisiting Post-Quantum Fiat-Shamir*. Cryptology ePrint Archive, Paper 2019/262. https://eprint.iacr.org/2019/262. 2019. URL: https://eprint.iacr.org/2019/262.

[98] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plancon. *Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General*. Cryptology ePrint Archive, Paper 2022/284. https://eprint.iacr.org/2022/284. 2022. URL: https://eprint.iacr.org/2022/284.

[99] Urmila Mahadev. *Classical Verification of Quantum Computations*. 2018. arXiv: 1804.01082 [quant-ph].

[100] Chris Marriott and John Watrous. *Quantum Arthur-Merlin Games*. 2005. arXiv: cs/0506068 [cs.CC].

[101] D. Micciancio and O. Regev. "Worst-case to average-case reductions based on Gaussian measures". In: *45th Annual IEEE Symposium on Foundations of Computer Science*. 2004, pp. 372–381. DOI: 10.1109/FOCS.2004.72.

[102] Daniele Micciancio and Chris Peikert. *Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller*. Cryptology ePrint Archive, Report 2011/501. https://eprint.iacr.org/2011/501. 2011.

[103] Daniele Micciancio and Oded Regev. “Worst-Case to Average-Case Reductions Based on Gaussian Measures”. In: *SIAM J. Comput.* 37.1 (2007), pp. 267–302. DOI: 10.1137/S0097539705447360. URL: https://doi.org/10.1137/S0097539705447360.

[104] Jörn Müller-Quade and Dominique Unruh. “Long-Term Security and Universal Composability”. In: *Theory of Cryptography*. Ed. by Salil P. Vadhan. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 41–60. ISBN: 978-3-540-70936-7.

[105] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. 10th. USA: Cambridge University Press, 2011. ISBN: 1107002176.

[106] Chris Peikert and Alon Rosen. “Efficient Collision-Resistant Hashing from Worst-Case Assumptions on Cyclic Lattices”. In: *Theory of Cryptography*. Ed. by Shai Halevi and Tal Rabin. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 145–166. ISBN: 978-3-540-32732-5.

[107] Chris Peikert and Sina Shiehian. “Privately constraining and programming PRFs, the LWE way”. In: *Public-Key Cryptography–PKC 2018: 21st IACR International Conference on Practice and Theory of Public-Key Cryptography, Rio de Janeiro, Brazil, March 25-29, 2018, Proceedings, Part II 21*. Springer. 2018, pp. 675–701.

[108] Daniele Perito and Gene Tsudik. *Secure Code Update for Embedded Devices via Proofs of Secure Erasure*. Cryptology ePrint Archive, Report 2010/217. https://ia.cr/2010/217. 2010.

[109] Alexander Poremba. “Quantum Proofs of Deletion for Learning with Errors”. In: *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*. Ed. by Yael Tauman Kalai. Vol. 251. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 90:1–90:14. DOI: 10.4230/LIPIcs.ITCS.2023.90. URL: https://doi.org/10.4230/LIPIcs.ITCS.2023.90.

[110] Roy Radian and Or Sattath. “Semi-Quantum Money”. In: *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*. ACM, Oct. 2019. DOI: 10.1145/3318041.3355462. URL: https://doi.org/10.1145%2F3318041.3355462.

[111] Oded Regev. “On Lattices, Learning with Errors, Random Linear Codes, and Cryptography”. In: *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*. STOC ’05. Baltimore, MD, USA: Association for Computing Machinery, 2005, pp. 84–93. ISBN: 1581139608. DOI: 10.1145/1060590.1060603. URL: https://doi.org/10.1145/1060590.1060603.

[112] Oded Regev. “On lattices, learning with errors, random linear codes, and cryptography”. In: *Journal of the ACM* 56.6 (2005), 34:1–34:40. ISSN: 0004-5411. DOI: 10.1145/1568318.1568324.

[113] R L Rivest, L Adleman, and M L Dertouzos. “On Data Banks and Privacy Homomorphisms”. In: *Foundations of Secure Computation, Academia Press* (1978), pp. 169–179.

[114] Ronald L. Rivest. “Can We Eliminate Certificate Revocation Lists?” In: *Proceedings Financial Cryptography ’98*. FC’98 (Anguilla, British West Indies, Feb. 23–25, 1998). Ed. by Rafael Hirschfeld. Vol. 1465. Lecture Notes in Computer Science. Springer, Feb. 1998, pp. 178–183. ISBN: 978-3-540-64951-9. DOI: 10.1007/BFb0055482.

[115] Bhaskar Roberts. *Toward Secure Quantum Money*. Princeton University Senior Thesis. http://arks.princeton.edu/ark:/88435/dsp01nc580q51r. 2019. URL: http://arks.princeton.edu/ark:/88435/dsp01nc580q51r.

[116] Bhaskar Roberts and Mark Zhandry. *Franchised Quantum Money*. Cryptology ePrint Archive, Paper 2021/1410. https://eprint.iacr.org/2021/1410. 2021. URL: https://eprint.iacr.org/2021/1410.

[117] Amit Sahai and Brent Waters. *How to Use Indistinguishability Obfuscation: Deniable Encryption, and More*. Cryptology ePrint Archive, Report 2013/454. https://eprint.iacr.org/2013/454. 2013.

[118] National Security Agency/Central Security Service. *Quantum Key Distribution (QKD) and Quantum Cryptography (QC)*. https://www.nsa.gov/Cybersecurity/Quantum-Key-Distribution-QKD-and-Quantum-Cryptography-QC/. [Online; accessed 05/01/2023].

[119] Omri Shmueli. *Public-Key Quantum Money with a Classical Bank*. Cryptology ePrint Archive, Paper 2021/1427. https://eprint.iacr.org/2021/1427. 2021. URL: https://eprint.iacr.org/2021/1427.

[120] Damien Stehlé et al. *Efficient Public Key Encryption Based on Ideal Lattices*. Cryptology ePrint Archive, Paper 2009/285. https://eprint.iacr.org/2009/285. 2009. URL: https://eprint.iacr.org/2009/285.

[121] S. Stubblebine. “Recent-secure authentication: enforcing revocation in distributed systems”. In: *2012 IEEE Symposium on Security and Privacy*. Los Alamitos, CA, USA: IEEE Computer Society, May 1995, p. 0224.

[122] Marco Tomamichel and Anthony Leverrier. “A largely self-contained and complete security proof for quantum key distribution”. In: *Quantum* 1 (July 2017), p. 14. ISSN: 2521-327X. DOI: 10.22331/q-2017-07-14-14. URL: https://doi.org/10.22331/q-2017-07-14-14.

[123] Marco Tomamichel et al. “A monogamy-of-entanglement game with applications to device-independent quantum cryptography”. In: *New Journal of Physics* 15.10 (2013), p. 103002.

[124] Matthias Troyer. *Towards Practical Quantum Advantage | Quantum Colloquium*. Youtube. 2021. URL: https://www.youtube.com/watch?v=WY3htdKUGsA&ab_channel=SimonsInstitute.

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

[126] Dominique Unruh. “Quantum proofs of knowledge”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2012, pp. 135–152.

[127] Dominique Unruh. "Revocable Quantum Timed-Release Encryption". In: *J. ACM* 62.6 (Dec. 2015). ISSN: 0004-5411. DOI: 10.1145/2817206. URL: https://doi.org/10.1145/2817206.

[128] Dominique Unruh. *Revocable quantum timed-release encryption*. Cryptology ePrint Archive, Paper 2013/606. https://eprint.iacr.org/2013/606. 2013. DOI: 10.1145/2817206. URL: https://eprint.iacr.org/2013/606.

[129] John Watrous. "Zero-Knowledge against Quantum Attacks". In: *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '06. Seattle, WA, USA: Association for Computing Machinery, 2006, pp. 296–305. ISBN: 1595931341. DOI: 10.1145/1132516.1132560. URL: https://doi.org/10.1145/1132516.1132560.

[130] Hoeteck Wee and Daniel Wichs. "Candidate obfuscation via oblivious LWE sampling". In: *Advances in Cryptology–EUROCRYPT 2021: 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17–21, 2021, Proceedings, Part III*. Springer. 2021, pp. 127–156.

[131] D. Wichs and G. Zirdelis. "Obfuscating Compute-and-Compare Programs under LWE". In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. 2017, pp. 600–611.

[132] Stephen Wiesner. "Conjugate Coding". In: *SIGACT News* 15.1 (Jan. 1983), pp. 78–88. ISSN: 0163-5700. DOI: 10.1145/1008908.1008920. URL: https://doi.org/10.1145/1008908.1008920.

[133] Mark M. Wilde. *Quantum Information Theory*. 1st. USA: Cambridge University Press, 2013. ISBN: 1107034256.

[134] A. Winter. "Coding theorem and strong converse for quantum channels". In: *IEEE Transactions on Information Theory* 45.7 (1999), pp. 2481–2485. ISSN: 0018-9448. DOI: 10.1109/18.796385. URL: http://dx.doi.org/10.1109/18.796385.

[135] W. K. Wootters and W. H. Zurek. "A single quantum cannot be cloned". In: *Nature* 299.5886 (Oct. 1982), pp. 802–803. DOI: 10.1038/299802a0. URL: http://dx.doi.org/10.1038/299802a0.

[136] Mark Zhandry. *Quantum Lightning Never Strikes the Same State Twice*. 2017. DOI: 10.48550/ARXIV.1711.02276. URL: https://arxiv.org/abs/1711.02276.

[137] Mark Zhandry. "Quantum Lightning Never Strikes the Same State Twice. Or: Quantum Money from Cryptographic Assumptions". In: *J. Cryptol.* 34.1 (Jan. 2021). ISSN: 0933-2790. DOI: 10.1007/s00145-020-09372-x. URL: https://doi.org/10.1007/s00145-020-09372-x.

[138] Mark Zhandry. *Schrödinger's Pirate: How To Trace a Quantum Decoder*. Cryptology ePrint Archive, Paper 2020/1191. https://eprint.iacr.org/2020/1191. 2020. URL: https://eprint.iacr.org/2020/1191.