

Department of Physics and Astronomy

University of Heidelberg

Master thesis

in Physics

submitted by

Alexander M. Poremba

born in Heidelberg, 14.12.1991

2017

Quantum Learning Algorithms and Post-Quantum Cryptography

This Master thesis has been carried out by Alexander Poremba

at

QMATH, Centre for the Mathematics of Quantum Theory,

University of Copenhagen

under the supervision of

Gorjan Alagic

and

under co-supervision of

Prof. Thomas Gasenzer

at the

University of Heidelberg.

Alexander M. Poremba

*University of Heidelberg, Department of Physics and Astronomy,
69124 Heidelberg, Germany.*

E-mail: alexander.poremba@gmail.com

ABSTRACT: Quantum algorithms have demonstrated promising speed-ups over classical algorithms in the context of computational learning theory - despite the presence of noise.

In this work, we give an overview of recent quantum speed-ups, revisit the Bernstein-Vazirani algorithm in a new learning problem extension over an arbitrary cyclic group and discuss applications in cryptography, such as the Learning with Errors problem. A defining advantage of quantum algorithms is parallelism, the ability to evaluate a function on a superposition of inputs. When solving a learning problem or attacking a cryptographic scheme, a quantum adversary can potentially exploit such parallelism and gain a substantial advantage over classical adversaries.

We turn to post-quantum cryptography and investigate new notions of security under non-adaptive quantum chosen-ciphertext attacks. In particular, we propose symmetric-key encryption schemes based on quantum-secure pseudorandom functions and permutations that fulfill our definitions. In order to prove security, we introduce a novel indistinguishability game, together with a blinding argument, and show that, in an oracle model, no quantum query algorithm making superposition queries can reliably distinguish between the class of Boolean functions that only differ at a single random location.

Finally, we discuss current progress in quantum computing technology and the implementation of quantum algorithms with a special focus on the ion-trap architecture. Moreover, we shed light on the relevance and effectiveness of common noise models adopted in computational learning theory.

ABSTRACT: Quantenalgorithmen sind vor allem im Bereich des maschinellen Lernens deutlich effizienter als konventionelle klassische Algorithmen - trotz Fehlerbehaftung.

In dieser Arbeit verschaffen wir einen Überblick über aktuell erzielte Erfolge, erweitern den bekannten Bernstein-Vazirani Algorithmus als ein Lernproblem über beliebigen zyklischen Gruppen und diskutieren potentielle Anwendungen in der Kryptographie, wie beispielsweise das Learning with Errors Problem. Ein wesentlicher Vorteil von Quantenalgorithmen besteht darin eine Funktion anhand von einer quantenmechanischen Überlagerung mehrerer Zustände auszuwerten. Aus diesem Grund hat ein Angreifer in Besitz eines Quantencomputers einen deutlichen Vorteil bei maschinellen Lernproblemen sowie bei Angriffen auf ein Kryptosystem gegenüber einem konventionellen Angreifer.

Wir beschäftigen uns außerdem mit der Post-Quanten-Kryptographie und untersuchen neue Sicherheitsstandards gegenüber nicht-adaptiven Angriffszenarien anhand von Quantenalgorithmen. Im Anschluss dazu führen wir symmetrische Kryptosysteme anhand pseudozufälliger Funktionen und Permutationen ein, welche sicher vor Angreifern in Besitz von Quantencomputern sind. Um die Sicherheit unserer Kryptosysteme zu beweisen, entwickeln wir ein neuartiges Sicherheitsspiel und zeigen damit, dass Quantenalgorithmen in einem Orakelmodell nicht zwischen Booleschen Funktionen unterscheiden können, die sich zufällig an einer Stelle unterscheiden.

Schließlich geben wir einen Überblick über den aktuellen Fortschritt in der Entwicklung von Quantencomputern, speziell über die Implementierung von Quantenalgorithmen in der Ionenfallen-Architektur. Außerdem untersuchen wir inwiefern übliche Fehlermodelle aus dem Bereich des maschinellen Lernens auch tatsächlich realistisch in der Praxis sind.

Contents

0 List of Abbreviations	4
1 Introduction	5
2 Main Results	9
3 Cryptography	12
3.1 Preliminaries	12
3.2 Symmetric-Key Cryptography	13
3.3 Security Notions	14
3.3.1 Computational Security	14
3.3.2 Computational Indistinguishability	14
3.3.3 Semantic Security	16
3.4 Pseudorandom Functions	17
3.5 Learning with Errors	18
3.5.1 Decision Learning with Errors	19
3.5.2 Symmetric-Key Constructions and Security	19
3.5.3 Separation Result	20
4 Quantum Computation	22
4.1 Formalism	22
4.2 Unitary Evolution	24
4.3 Quantum Measurement	25
4.4 Universal Quantum Gates	26
4.5 The No-Cloning Theorem	28
4.6 The Quantum Circuit Model	28
4.7 Quantum Parallelism	29
4.8 Decoherence	31
4.8.1 Quantum Noise Models	31
4.8.2 Independent Noise Models	32
4.9 Error Correcting Codes	33
4.10 Quantum Oracles	35
4.10.1 Membership Oracles	36
4.10.2 Example Oracles	36
5 Quantum Algorithms	38
5.1 Deutsch-Josza Algorithm	39
5.2 Bernstein-Vazirani Algorithm	40
5.3 Simon's Algorithm	42

6 The Quantum Fourier Transform	44
6.1 The Quantum Fourier Transform over Finite Abelian Groups	45
6.2 Efficient Circuit Implementations	48
7 Quantum Learning Algorithms	50
7.1 Computational Learning Theory	50
7.1.1 Exact Learning	51
7.1.2 PAC Learning	51
7.2 Learning Parity With Noise	53
7.3 Extended Bernstein-Vazirani Algorithm	56
7.4 Learning With Errors	58
8 Blinding of Quantum Algorithms	59
8.1 Blinding Lemma	59
8.2 Relabeling Games	61
9 Post-Quantum Cryptography	63
9.1 Security Under Non-adaptive Quantum Chosen-Ciphertext Attacks	63
9.1.1 Indistinguishability	64
9.1.2 Decisional Indistinguishability	64
9.1.3 Semantic Security	65
9.1.4 Equivalence of Indistinguishability and Semantic Security	65
9.2 Quantum-secure Pseudorandom Functions and Permutations	66
9.3 Secure Constructions	67
10 The Physical Realization of Quantum Computation	75
10.1 DiVincenzo Criteria	75
10.2 Ion-Trap Implementation	77
10.2.1 Hyperfine Structure	77
10.2.2 Experimental Setup	78
10.2.3 The Hamiltonian	81
10.2.4 Single-Qubit Gates	84
10.2.5 Two-Qubit Gates	86
10.2.6 Quantum Algorithms with Trapped Ions	88
10.2.7 Decoherence and Sources of Error	91
10.3 Performance Comparison of Quantum Computing Architectures	93
11 Conclusion	93
12 Open Problems	94
A Supplementary Material	95

Acknowledgments

I want to thank Gorjan Alagic for his support and mentorship, and for inspiring me to pursue a thesis in cryptography. I also want to thank Thomas Gasenzer for making this thesis possible.

For my parents.

0 List of Abbreviations

\perp *reject symbol*

CPA *chosen-plaintext attack*

CPTP *completely positive and trace preserving*

CCA1 *non-adaptive chosen-ciphertext attack*

CCA2 *adaptive chosen-ciphertext attack*

DecIND *decisionally indistinguishable encryptions*

DecIND-QCPA *decisional indistinguishability under quantum chosen-plaintext attack*

DecIND-QCCA1 *decisional indistinguishability under non-adaptive quantum chosen-ciphertext attack*

DecLWE *Decision Learning with Errors*

IND *indistinguishable encryptions*

IND-CPA *indistinguishable encryptions under chosen-plaintext attack*

IND-CCA1 *indistinguishable encryptions under non-adaptive chosen-ciphertext attack*

IND-CCA2 *indistinguishable encryptions under adaptive chosen-ciphertext attack*

IND-QCPA *indistinguishable encryptions under quantum chosen-plaintext attack*

IND-QCCA1 *indistinguishable encryptions under non-adaptive quantum chosen-ciphertext*

IND-QCCA2 *indistinguishable encryptions under adaptive quantum chosen-ciphertext attack*

LWE *Learning with Errors*

LPN *Learning Parity with Noise*

NP *nondeterministic polynomial time*

P *polynomial time*

PAC *probably approximately correct*

PPT *probabilistic polynomial time*

POVM *positive operator valued measure*

PRF *pseudorandom function*

PRP *pseudorandom permutation*

QCPA *quantum chosen-plaintext attack*

QCCA1 *non-adaptive quantum chosen-ciphertext attack*

QCCA2 *adaptive quantum chosen-ciphertext attack*

QFT *quantum Fourier transform*

QPRF *quantum-secure pseudorandom function*

QPRP *quantum-secure pseudorandom permutation*

QPT *quantum polynomial time*

SKES *symmetric-key encryption scheme*

SEM *semantic security*

SEM-CCA1 *semantic security under non-adaptive chosen-ciphertext attack*

SEM-CCA2 *semantic security under adaptive chosen-ciphertext attack*

SEM-QCCA1 *semantic security under non-adaptive quantum chosen-ciphertext attack*

1 Introduction

Most of our present day communication takes place on the internet and produces enormous amounts of personal data. Whereas traditional notions of security are mainly concerned with electronic mail or bank transfers, today's security needs have since expanded to many unexpected areas such as smartcards, medical devices or modern cars. Cryptography, understood as the science of secure communication, is becoming increasingly relevant for our safety in the modern world. For many years, popular cryptographic protocols such as RSA, the Diffie-Hellman key-exchange (D-H) or elliptic curve cryptography (ECC), have served greatly as building blocks to establish secure communication, despite lower costs and ever increasing computational power on the markets. It was Peter Shor's 1994 breakthrough discovery of efficient quantum algorithms for the factoring of integers and the computation of discrete logarithms [Sho94] that truly drew the attention towards the field of quantum computation and its potential impact on cryptography. Many of the protocols still in use today, such as RSA, D-H or ECC, are completely broken by attackers in possession of quantum computers running Shor's algorithm. This discovery is oftentimes regarded as the beginning of a new race towards *post-quantum cryptography*, a security standard for secure classical communication, even in the presence of quantum computers [BL17]. While quantum cryptography in itself has provided us with entirely new forms of communication, such as quantum key distribution [NC10], it is reasonable to assume that some form of classical communication will nevertheless continue to exist for years to come. Even as reliably fault-tolerant quantum computers have yet to be built, the cryptographic community has nevertheless started shifting towards a new direction in which the feasibility of classical cryptography in a quantum world presents us with an important challenge.

A common approach in cryptography is the integration of hard computational problems towards the implementation of secure communication. Consider, for example, the well known RSA protocol whose security is based on the fact that factoring large integers is believed to be computationally intractable. Ever since the discovery of Shor's algorithm, the search towards computational hardness in a quantum world has dominated the cryptographic community. Since 2005, the *Learning with Errors* (LWE) problem [Reg05] has gained the status of a promising cryptographic basis of hardness, in particular in a post-quantum setting. The central promise of the LWE problem lies in a reduction in which it is shown to be as hard as worst-case lattice problems [Reg09], a class of computational problems believed to be hard for more than two decades. Consequently, it is tempting to build cryptographic constructions on the basis of the LWE problem and achieve security under the assumption that worst-case lattice problems remain hard for quantum computers. Apart from being a candidate for security against quantum computers, companies like *Google* and *IBM* have also shown interest in variants of LWE due to its promise for light-weight implementation. As of today, the security of lattice-based cryptography against quantum computers remains one of the key areas of modern research in cryptography.¹

¹For an excellent review on modern cryptography in the age of quantum computers, we refer to a popular science article in a 2015 issue of Quanta Magazine: www.quantamagazine.org/quantum-secure-cryptography-crosses-red-line-20150908/

In a nutshell, given an integer n , a modulus q and secret string $s \in (\mathbb{Z}/q\mathbb{Z})^n$, the **LWE** problem can be stated as follows:

Recover a secret string s given a set of noisy linear equations on s .

In particular, let us consider the case in which the length of the string is $n = 4$, the modulus is given by $q = 23$ and (with probability less than $1/2$) each equation is of small additive error ± 1 . The goal is now to find s based on the following set of noisy linear equations, where each sample of coefficients on s is chosen uniformly at random:

$$\begin{aligned}
11s_1 + 2s_2 + 13s_3 + 19s_4 &\approx 8 \pmod{23} \\
14s_1 + 6s_2 + 19s_3 + s_4 &\approx 5 \pmod{23} \\
3s_1 + 15s_2 + 4s_3 + 2s_4 &\approx 0 \pmod{23} \\
4s_1 + 6s_2 + 20s_3 + 15s_4 &\approx 11 \pmod{23} \\
7s_1 + 18s_2 + 8s_3 + 9s_4 &\approx 21 \pmod{23} \\
8s_1 + 5s_2 + 17s_3 + 12s_4 &\approx 10 \pmod{23} \\
&\vdots \\
16s_1 + s_2 + 11s_3 + 22s_4 &\approx 14 \pmod{23}
\end{aligned}$$

If q is prime, the integers modulo q form a finite field under addition and multiplication, hence, given enough samples on s , there exists a unique solution to the problem. In our case, the hidden string to be determined is $s = (12, 0, 7, 2)$. If not for the error, the secret string can be recovered in polynomial time $O(n^3)$ using Gaussian elimination after observing n linear independent equations, where n denotes the length of the string. Let us also note that the probability of acquiring n linear independent equations on s after only observing n sample queries is easily shown to be both greater than a constant and independent on n .

The difficulty in decoding noisy linear equations lies in the fact that the errors propagate during the computation, hence amplify the uncertainty and ultimately lead to no information on the actual secret string. As the best known algorithm for the **LWE** problem runs in time $O(2^n)$ [BKW03], the problem is believed to be asymptotically intractible for classical computers. Moreover, due to the reduction in [Reg05], any breakthrough in **LWE** would also most likely imply an algorithm for lattice-based problems.

In an earlier problem, Bernstein and Vazirani [BV93] considered the task of determining a hidden string from inner product of bit strings in a setting where an algorithm is granted input access to evaluations of the function (here \oplus denotes addition modulo 2):

Bernstein-Vazirani Problem:

Recover a string $s \in \{0, 1\}^n$ by making input queries to a Boolean function given by $f_s : \{0, 1\}^n \rightarrow \{0, 1\}$, where

$$f_s(x) = s_1 \cdot x_1 \oplus \dots \oplus s_n \cdot x_n = \langle s, x \rangle \bmod 2.$$

This problem features a curious resemblance to a variant of the LWE problem in which the modulus is given by $q = 2$, the algorithm is free to choose all inputs (instead of receiving samples uniformly at random) and where the noise is absent from all evaluations of the function. In the classical query setting, we observe that a single query to the function can only reveal as much as one bit of information about the secret string s . In fact, this can easily be done by considering queries on strings $e_i = (0, \dots, 1, \dots, 0)$, where the i -th index is 1 and e_i is 0 everywhere else. Any algorithm performing such queries thus achieves an overall query complexity of $O(n)$ when determining the secret, as each query reveals a single bit:

$$f_s(e_i) = \langle s, e_i \rangle \bmod 2 = s_i, \quad (1.1)$$

so that s is fully determined after a total of n queries to the function. Therefore, it is tempting to approach the LWE problem in this simplified model.

In this thesis, we consider the Bernstein-Vazirani problem in a setting in which an algorithm is given quantum access to the function, hence is able to exploit quantum parallelism and to evaluate the inner product simultaneously on a superposition of inputs. More formally, the algorithm can evaluate f_s through a quantum operation, a black box whose inner workings regarding the computation of the function are unknown to the algorithm. In particular, we introduce the notion of an *oracle*, a quantum operation \mathcal{O}_{f_s} that allows for the reversible evaluation of a function f upon a set of inputs as follows:

$$\mathcal{O}_{f_s} : \sum_{x, y \in \{0, 1\}^n} \alpha_{x, y} |x\rangle |y\rangle \longrightarrow \sum_{x, y \in \{0, 1\}^n} \alpha_{x, y} |x\rangle |y \oplus f_s(x)\rangle. \quad (1.2)$$

Remarkably, as Bernstein and Vazirani [BV93] showed, only a single oracle query to the function as in (1.2) is sufficient to determine the secret string. We generalize this model to a cyclic group $\mathbb{Z}/q\mathbb{Z}$ of arbitrary integers q in a new learning problem extension of the Bernstein-Vazirani algorithm and discuss its speed-up over classical algorithms. Cross et al. [CSS14] have recently demonstrated a robustness of quantum learning for certain classes of noise in which samples are also likely to be corrupted. While this setting is known to cause most learning problems intractable for classical algorithms, the analogue using quantum samples remains easy. This fact recently allowed Grilo et al. [GK17] to independently find an efficient quantum learning algorithm for LWE, a special variant of our proposed extended Bernstein-Vazirani algorithm in which q is prime. While this algorithm does not solve the LWE problem in its original formulation using classical samples, it does however suggest further caution when allowing access to quantum samples in any cryptographic application. Nevertheless, not even a quantum computer receiving classical LWE samples,

i.e. classical strings of noisy linear equations, seems to be able to challenge the hardness of LWE [Reg09]. For this reason, LWE is believed to be an excellent basis of hardness in post-quantum cryptography.

While quantum superposition access is regularly shown to be a powerful model, it also possesses limitations. Our goal in this work is also to find such limitations in order to provide quantum-secure encryption schemes, even in a setting in which an attacker has quantum access to the encryption procedure. An essential building block for the construction of secure cryptographic schemes is found in so-called pseudorandom functions, a family of keyed functions that seem indistinguishable from random functions to any adversary with limited computational resources. In fact, recent breakthroughs in quantum cryptography allow for quantum-secure pseudorandom functions that are secure, even if an adversary is given the ability to evaluate the function using quantum superpositions. Remarkably, as shown by Zhandry in 2012, such constructions can be built using the classical sample hardness of LWE in the quantum world [Zha12]:

If LWE with classical samples is hard for quantum computers, then there exist quantum-secure pseudorandom functions.

As parallelism remains one of the key features of quantum algorithms, the goal is to exploit the nature of complex-valued amplitudes of quantum states and cause them to interfere around the desired outputs through the use of quantum operations. Only then, a final measurement of the state collapses the superposition into the desired outcome with high probability. The following fact guarantees that quantum parallelism can be achieved for all efficiently computable functions [NC10]:

Any classical efficiently computable function has an efficient circuit description, hence can also be implemented efficiently using a quantum computer. Moreover, the quantum circuit for the function consists entirely of unitary gates and can thus be evaluated on a superposition of inputs due to the linearity of quantum mechanics.

A fundamental question arises immediately. Just how powerful is knowledge represented in a quantum superposition evaluating a function on all of its inputs? This thesis is concerned with both the limitations and exploitation of quantum parallelism in the context of modern cryptography.

2 Main Results

Let us now give an overview of the main contents provided in this thesis.

In CHAPTER 3, we review selected topics modern cryptography required for the proposed constructions in this thesis. We introduce the concept of *symmetric-key encryption schemes* (SKES), a setting in which two agents, say Alice and Bob, share a matching secret key prior to their communication. In order to encrypt messages, Alice first runs an encryption algorithm that requires the use of her key and later sends the resulting ciphertext over to Bob. Since Bob knows about the secret key, he can run a decryption algorithm upon Alice’s ciphertext and decode the message. In order to prove the security of symmetric-key encryption schemes, we introduce several relevant notions of security. Moreover, we define the concept of *pseudorandom functions*, a crucial building block in symmetric-key cryptography that allows for constructions of SKES of precisely such security. Moreover, we quantify limited computational power by introducing the notion of *efficient* adversaries who run algorithms with at most polynomial running time with regard to some security parameter relevant to the underlying cryptographic scheme. Finally, we define the LWE problem rigorously and discuss its applications in cryptography.

In CHAPTER 4, we present the most important developments in the theory of quantum computation to date. To this end, we introduce the concept of qubits, unitary quantum operations and the quantum circuit model. We present a universal set of quantum gates that enables a quantum computer to approximately perform any quantum operation. Furthermore, we discuss the *no-cloning theorem*, a fundamental feature of quantum mechanics that forbids the copying of quantum information. This fact highlights one of the defining aspects of quantum cryptography and provides a crucial assumption for the majority of the work in this thesis. Most importantly, we give examples of quantum parallelism and show how to prepare a quantum state that evaluates a given function simultaneously over the range of its inputs. In this context, we introduce the concept of quantum oracles, essentially a quantum gate that acts as a black box and grants an algorithm input access to a given function. Finally, we turn to noise and decoherence in quantum computing architectures and give examples of elementary error correcting codes.

In CHAPTER 5, we review several of the well known quantum algorithms that solve certain computational tasks faster than any known classical algorithm and provide the foundation for the algorithms of the later chapters. In particular, we introduce the *Deutsch-Josza algorithm*, the earliest quantum speed-up ever to be found in a black box model. Moreover, we introduce the *Bernstein-Vazirani algorithm* as the original predecessor of the *extended Bernstein-Vazirani algorithm*. Finally, we present *Simon’s period finding algorithm*, a highly useful tool to attack *block-ciphers* in symmetric-key cryptography.

In CHAPTER 6, we introduce the quantum Fourier transform (QFT) over arbitrary finite abelian groups as a fundamental operation adopted in the majority of all the algorithms discussed in this thesis. The Fourier transform is particularly useful in exploiting the symmetries of a given problem and allows us to generalize the Bernstein-Vazirani algorithm over arbitrary cyclic groups. Finally, we discuss efficient quantum circuit implementations that compute the QFT.

In CHAPTER 7, we introduce useful language from *computational learning theory* in which we frame the main algorithms in this thesis. We consider a setting in which a learner (an algorithm) is asking for samples from a black box oracle whose inner workings are unknown. The goal of the learner is to determine a *concept*, such as a secret boolean function, based on the information that is being presented by the samples. As samples may be subjected to noise, errors are likely to get amplified, resulting in oftentimes highly non-trivial tasks that are computationally intractable for classical computers. In particular, we consider the *Learning Parity with Noise* (LPN) problem, an early predecessor of the LWE problem, as an instance of a computational learning problem. Once we define the analogous learning problem in a setting in which the oracle is providing quantum samples, we investigate how these computational tasks become easy for quantum computers. We approach a quantum LWE analogue by first proposing a new generalization of the Bernstein-Vazirani algorithm over an arbitrary cyclic group and show that the secret string can be determined with high probability.

Algorithm 6. *There exists a quantum algorithm for the extended Bernstein-Vazirani problem that can be amplified towards a success probability of $1 - \delta$ by requesting $O(\log 1/\delta)$ many samples independently of n , whereas the classical query complexity is given by $\Omega(n)$.*

In addition, we compare our results to an independent 2017 proposal by Grilo and Kerenidis that proves that, in the quantum oracle setting, the extended Bernstein-Vazirani algorithm (in the special case where q is prime) solves the LWE problem given enough quantum samples.

In CHAPTER 8, we take a turn towards studying the limitations of quantum algorithms in order to find secure constructions for post-quantum cryptography. While the previous chapter focused on quantum speed-ups at solving learning problems by means of superposition samples, this chapter investigates the limitations of quantum algorithms instead. We discuss the effects of blinding of quantum algorithms, a setting in which we modify the function to which the algorithm is given oracle access to at a single location and study its subsequent output states. In particular, we prove a *blinding lemma* that states that any quantum query algorithm produces output states that remain negligibly close in expected trace distance, irrespective of modification at a single random location. Finally, we introduce a new indistinguishability game, the *RelabelingGame*, a setting in which a distinguisher has quantum oracle access to a known function and a challenger randomly modifies the output of the function at a single location halfway through the game. The goal of the distinguisher is thus to determine whether such modification occurred after an initial query phase. Using a blinding argument, we prove:

Proposition 8.3. *Let $f : \{0,1\}^n \rightarrow \{0,1\}^m$ be a function. Then any quantum polynomial time algorithm \mathcal{D} making oracle queries to \mathcal{O}_f wins the RelabelingGame with at most negligible probability $\frac{1}{2} + \epsilon(n)$.*

In CHAPTER 9, we extend the notions of classical indistinguishability to a quantum world. We make use of the blinding argument and propose secure constructions under a quantum chosen-ciphertext attack. In this scenario, a quantum adversary exercises control over the functionality of the scheme and is able to influence an honest party into quantumly generating ciphertexts, as well as decrypting ciphertexts of the adversary's choice for some period in time. We introduce several notions of security, such as decisionally indistinguishable encryptions and semantic security, and prove that our proposed constructions satisfy our definitions. Finally, we give two secure constructions based on either quantum-secure pseudorandom functions, or quantum-secure pseudorandom permutations.

In CHAPTER 10, we discuss state-of-the-art quantum computing technology with a particular focus on the ion-trap architecture. We give a detailed introduction to how qubits are realized in a physical system and how quantum gates can be performed through the use of lasers. Furthermore, we discuss sources of noise and decoherence in physical systems in order to investigate the effectiveness of noise models from the previous chapters. Finally, we discuss the performance of recent implementations of quantum algorithms discussed in this thesis.

3 Cryptography

The history of cryptography dates back to over two millenia. Ever since the birth of civilization and the invention of writing, people required ways of transmitting secret messages using *ciphers*, intended to be read only by the receiver and yet difficult to decode for others. Since the 1970s, cryptography amounted to a well-established scientific discipline by henceforth adopting a rigorous mathematical foundation. This crucial change marks the beginning of *modern cryptography*. Many of the popular encryption schemes still in use today, such as the RSA encryption scheme, were already developed in the early years of modern cryptography. Typically, it is the hardness of certain computational problems that serves as a foundation for security. For example, as in the case of RSA, the security of the encryption scheme is related to the hardness of factoring large integers. In other words, we believe a scheme is secure, if no efficient adversary with limited computational resources is capable of breaking the scheme. Peter Shor's discovery of an efficient quantum algorithm for the factoring of integers marked the beginning of an entirely new era of cryptography, a so-called *post-quantum cryptography*. It is from here on, that the search for quantum-secure cryptography began. In the following sections, we provide an overview of selected topics in modern cryptography required for the main results in this thesis.

3.1 Preliminaries

Let us first introduce some necessary notation and formalism from theoretical computer science and cryptography. For additional reading, we refer to [KL15].

For bit strings $x \in \{0, 1\}^n$ of arbitrary length $n = |x|$, we associate a product space $\{0, 1\}^*$ containing all such strings of finite length. A function $\epsilon : \mathbb{N} \rightarrow \mathbb{R}$ is called *negligible* if, for every polynomial p , there exists an integer N such that for all $n > N$, it holds that: $\epsilon(n) < \frac{1}{p(n)}$. Typically, we adopt negligible functions in the context of a success probability that decreases to an inverse-superpolynomial rate, hence cannot be amplified to a constant by a polynomial amount of repetitions. An algorithm is a sequence of (possibly nondeterministic) operations that terminates after a finite amount of steps upon any given input, say $x \in \{0, 1\}^*$. We say an algorithm is *efficient* if it has polynomial running time with respect to a size parameter of a given computational problem, i.e. if there exists a polynomial $p(x)$ such that, for any input $x \in \{0, 1\}^*$, the computation of $A(x)$ terminates after at most $p(|x|)$ steps. A probabilistic polynomial time (PPT) algorithm is a procedure with an additional random tape (such as a random number generator) that results in efficient, yet possibly nondeterministic, computations. We adopt the popular unary convention of representing the seed of efficient randomized algorithms by $1^n = 11\dots 1$, highlighting a polynomial dependence with respect to the length of the input, contrary to a polylog dependence in the general case where $\lceil \log_2(n) \rceil$ bits are needed to specify the length of a given input (here, $\lceil \cdot \rceil$ denotes the ceiling function). With $x \leftarrow X$, we denote a procedure an outcome x is sampled uniformly at random from a finite set X . If D is a probability distribution, we denote the sampling of an outcome according to D by using the notation $x \leftarrow D$. Upon finite sets X and Y , we define the corresponding (finite) set of all possible functions from X to Y as $\{\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}\}$. An *oracle* is a black box machine \mathcal{O} that assists a given algorithm with a par-

ticular computational task at unit cost, for example an evaluation of an unknown function upon a given input or the sampling from an unknown probability distribution. Typically, if \mathcal{A} is an algorithm, we denote oracle access to \mathcal{O} using the notation $\mathcal{A}^{\mathcal{O}}$. Finally, throughout this thesis, we employ the usual asymptotic O -notation denoting an upper bound, where for a given function $g(n)$, we define $O(g(n)) = \{f(n) : \exists c \in \mathbb{R}, \exists n \in \mathbb{N} \text{ such that } 0 \leq f(n) \leq c g(n), \forall n \geq n_0\}$. Similarly, we denote an asymptotic lower-bound $\Omega(g(n))$ as the set of functions $\Omega(g(n)) = \{f(n) : \exists c \in \mathbb{R}, \exists n \in \mathbb{N} \text{ such that } 0 \leq c g(n) \leq f(n), \forall n \geq n_0\}$.

3.2 Symmetric-Key Cryptography

Symmetric-key cryptography concerns the scenario in which two agents, say Alice and Bob, share a mutual secret key k prior to their communication and want to send messages to each other. In order to encrypt messages, Alice first chooses a message m and runs an encryption algorithm $\text{Enc}_k(m)$ that requires the use of her key and later sends the resulting ciphertext c over to Bob. Since Bob knows about the secret key, he can run a decryption algorithm $\text{Dec}_k(c)$ upon Alice's ciphertext and decode the message. In general, we consider randomized encryption in order to avoid *replay attacks*, while only requiring decryption to be deterministic.

Definition 3.1. *A symmetric-key encryption scheme (SKES) is a triple of PPT algorithms $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ on a finite key space \mathcal{K} , message space \mathcal{M} and ciphertext space \mathcal{C} , where $\text{KeyGen} : \mathbb{N} \rightarrow \mathcal{K}$, $\text{Enc} : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{C}$, $\text{Dec} : \mathcal{K} \times \mathcal{C} \rightarrow \mathcal{M}$ and, for a security parameter n , we require:*

1. (key generation) KeyGen : on input 1^n , generate a key $k \leftarrow \text{KeyGen}(1^n)$;
2. (encryption) Enc_k : on message $m \in \mathcal{M}$, output a ciphertext $\text{Enc}_k(m)$;
3. (decryption) Dec_k : on cipher $c \in \mathcal{C}$, output a message $\text{Dec}_k(c)$;
4. (correctness) $(\text{Dec}_k \circ \text{Enc}_k)(m) = m$.

In order for communication under a given symmetric-key encryption scheme to be secure against eavesdroppers, we require that, without knowledge of the secret key, any ciphertext must look sufficiently random and reveal little to no information about the actual message.

In the next section, we provide several widely used notions of security for symmetric-key encryption. For further reading, we refer to [KL15].

3.3 Security Notions

3.3.1 Computational Security

Due to the well known P-NP problem, i.e. the seeming impossibility of finding efficient algorithms for certain hard computational problems, and the fact that we consider adversaries who operate probabilistically, an important notion of security is provided by *computational security* based on the following principle:

A cipher must be practically secure (if not mathematically secure) against adversaries with limited computational resources.

This brings us to the following definition of computational security:

Definition 3.2 (Computational Security). *A scheme Π is computationally (or asymptotically) secure if every PPT adversary succeeds at breaking Π with at most negligible probability with respect to the security parameter of Π .*

Since a negligible success probability is smaller than the inverse of any polynomial, no efficient algorithm is capable of amplifying the success probability, i.e. capable of breaking the encryption scheme by sheer repetition. Therefore, we regard any algorithm that breaks a particular scheme with at most negligible probability as not significant.

3.3.2 Computational Indistinguishability

Another important notion of security for a given symmetric-key encryption scheme is *indistinguishability of encryptions*, in particular under a *chosen-plaintext attack*. In this model, an adversary has partial control over the encryption procedure and can generate encryptions of arbitrary messages. This attack corresponds to a scenario in which an attacker is able to influence an honest party into generating ciphertexts of the adversaries choice, thus potentially resulting in an advantage at decoding other ciphers of interest. In the following, we specify this model in a security game between an adversary and a challenger:

Definition 3.3 (IND-CPA).

Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme and consider the INDGame between a PPT adversary and challenger, defined as follows:

1. (initial phase) *the challenger chooses a key $k \leftarrow \text{KeyGen}(1^n)$ and bit $b \leftarrow \{0, 1\}$;*
2. (pre-challenge phase) *as part of a learning phase, the adversary is given access to an encryption oracle Enc_k in order to generate encryptions. Upon each choice of message m , the adversary receives a ciphertext $c \leftarrow \text{Enc}_k(m)$. Finally, the adversary chooses two messages m_0 and m_1 , and sends them to the challenger.*
3. (challenge phase) *the challenger replies with $\text{Enc}_k(m_b)$ and the adversary continues to have oracle access to Enc_k ;*
4. (resolution) *the adversary outputs a bit b' and wins the game if $b' = b$.*

We say that Π has indistinguishable encryptions under a chosen-plaintext attack (or is IND-CPA-secure) if, for every PPT \mathcal{A} , there exists a negligible function $\epsilon(n)$ such that: $\Pr[\mathcal{A} \text{ wins INDGame}] \leq 1/2 + \epsilon(n)$.

An even stronger notion of security for a given symmetric-key encryption scheme is *security under chosen-ciphertext attacks*. In this variant of the INDGame, an adversary not only exercises control over the encryption scheme as before, but can also *non-adaptively* decrypt messages unrelated to a ciphertext of interest (as highlighted in the pre-challenge and challenge phase). Therefore, such an attack corresponds to a scenario in which an attacker is able to exercise control over an honest party into generating ciphertexts, as well as decrypting ciphertexts of the adversary's choice for some period in time. In the following, we specify this model in another security game between an adversary and a challenger:

Definition 3.4 (IND-CCA1).

Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme and consider the INDGame between a PPT adversary and challenger, defined as follows:

1. (initial phase) the challenger chooses a key $k \leftarrow \text{KeyGen}(1^n)$ and bit $b \leftarrow \{0, 1\}$;
2. (pre-challenge phase) as part of a learning phase, the adversary is given access to both an encryption oracle Enc_k and decryption oracle Dec_k . Upon each choice of message m , the adversary receives a ciphertext $\text{Enc}_k(m)$ and, upon each ciphertext c , the adversary receives a plaintext $\text{Dec}_k(c)$. Finally, the adversary chooses two messages m_0 and m_1 , and sends them to the challenger.
3. (challenge phase) the challenger replies with $\text{Enc}_k(m_b)$ and the adversary continues to have oracle access to Enc_k only;
4. (resolution) the adversary outputs a bit b' and wins the game if $b' = b$.

We say that Π has indistinguishable encryptions under a chosen-ciphertext attack (or is IND-CCA1-secure) if, for every PPT \mathcal{A} , there exists a negligible function $\epsilon(n)$ such that: $\Pr[\mathcal{A} \text{ wins INDGame}] \leq 1/2 + \epsilon(n)$.

Finally, we can additionally extend the previous notion of IND-CCA1 security by also granting the adversary *adaptive* decryption access after the challenge phase. This model corresponds to IND-CCA2 security, a variant in which the adversary exercises full control over the encryption scheme, both before and after the challenge phase. Remarkably, there exist classical symmetric-key encryption schemes that satisfy each of the security definitions provided in this chapter. A major contribution of this thesis is to provide constructions that satisfy these notions, even in a setting in which the adversary is granted quantum superposition access, again both to the encryption and decryption procedure. In the next section, we introduce important tools to realize such cryptographic schemes.

3.3.3 Semantic Security

In semantic security, the challenge phase corresponds to choosing a *challenge template* instead of a pair of messages. Contrary to the **INDGame**, the intuition for this security game is that the adversary seeks to compute something meaningful about the message of interest during the challenge phase. Thus, we consider challenge templates consisting of a triple of classical circuits (\mathbf{Samp}, h, f) , where \mathbf{Samp} outputs plaintexts from some distribution $\mathcal{D}_{\mathbf{Samp}}$, and h and f are functions over messages $m \leftarrow \mathbf{Samp}$. Upon receiving an encryption of a randomly sampled message m according to \mathbf{Samp} , the goal of the adversary is to output some new information $f(m)$, given some side information $h(m)$ on the message. In providing an adversary with a **CCA1** learning phase, we can consider the following notion of security.

Definition 3.5 (SEM-CCA1). *Let $\Pi = (\mathbf{KeyGen}, \mathbf{Enc}, \mathbf{Dec})$ be an encryption scheme, and consider the experiment **SEMGame** with a PPT \mathcal{A} , defined as follows.*

1. (initial phase) *A key $k \leftarrow \mathbf{KeyGen}(1^n)$ and bit $b \xleftarrow{\$} \{0, 1\}$ are generated;*
2. (pre-challenge phase) *\mathcal{A} receives access to oracles \mathbf{Enc}_k and \mathbf{Dec}_k , then outputs a challenge template consisting of (\mathbf{Samp}, h, f) ;*
3. (challenge phase) *A plaintext $m \leftarrow \mathbf{Samp}$ is generated; \mathcal{A} receives $h(m)$ and an oracle for \mathbf{Enc}_k only; if $b = 1$, \mathcal{A} also receives $\mathbf{Enc}_k(m)$.*
4. (resolution) *\mathcal{A} outputs a string s , and wins if $s = f(m)$.*

We say Π is semantically secure under a non-adaptive chosen-ciphertext attack (or is **SEM-CCA1**) if, for every PPT \mathcal{A} , there exists a PPT \mathcal{S} such that the challenge templates output by \mathcal{A} and \mathcal{S} are identically distributed, and there exists a negligible function $\epsilon(n)$ such that:

$$\left| \Pr_{k \xleftarrow{\$} \mathcal{K}} [\mathcal{A}(1^n, \mathbf{Enc}_k(m), h(m)) = f(m)] - \Pr[\mathcal{S}(1^n, |m|, h(m)) = f(m)] \right| \leq \epsilon(n),$$

where, in both cases, the probability is taken over plaintexts $m \leftarrow \mathbf{Samp}$.

Fortunately, as shown in [KL15], semantic security and indistinguishability are equivalent notions of security, in particular under non-adaptive chosen-ciphertext attacks.

Theorem 3.6. *Let $\Pi = (\mathbf{KeyGen}, \mathbf{Enc}, \mathbf{Dec})$ be a symmetric-key encryption scheme. Then, Π is **IND-CCA1**-secure if and only if Π is **SEM-CCA1**-secure.*

In Chapter 9, we introduce variants under quantum chosen-ciphertext attacks and prove the equivalence of both definitions. While semantic security is a much more intuitive notion of security, it is oftentimes much harder to prove security in practice. Therefore, it is convenient to provide security proofs under indistinguishable encryptions and refer to such an equivalence for a more natural notion of security.

3.4 Pseudorandom Functions

In this section, we turn to pseudorandom functions, a popular building block in symmetric-key cryptography. Historically, the first instance of provably-secure pseudorandom functions was proposed in the Goldreich, Goldwasser and Micali construction [GGM86] using pseudorandom generators, which in turn rely on the existence of one-way functions. The effectiveness of pseudorandom functions lies in the property of seeming indistinguishable from a perfectly random function to any efficient distinguisher with limited computational power. The security properties of pseudorandom functions are perhaps best explained in an indistinguishability game between a distinguisher (a PPT algorithm) and a challenger. Upon the start of the game, the challenger chooses a random bit b whose outcome determines whether the game is being played with a perfectly random function (sampled uniformly at random from the finite set of all possible functions over given finite domain and range) of the challengers choice, or a pseudorandom function for a freshly generated key. Next, the challenger presents the distinguisher with an oracle for the given function who is then free to evaluate the function upon arbitrary inputs. Finally, the distinguisher wins by outputting a bit $b' = b$. Since the distinguisher is assumed to have limited computational resources, thus essentially running a PPT algorithm, the claim of pseudorandomness is that the outputs will look sufficiently random. Therefore, the probability that the distinguisher makes a decision in a game against a pseudorandom function and outputs a bit, say $b' = 1$, is negligibly close to a game in which the distinguisher is playing against a perfectly random function. We formalize this observation in the following definition:

Definition 3.7. (Pseudorandom Function)

Let f be a keyed function $f : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ on a key-space \mathcal{K} , a domain \mathcal{X} and a range \mathcal{Y} . We say $\text{PRF} = \{f_k\}_{k \in \mathcal{K}}$ is a family of pseudorandom functions if, for every choice of key k and for all PPT distinguishers \mathcal{D} , there exists a negligible function $\epsilon(n)$ such that:

$$\left| \Pr_{k \xleftarrow{\$} \mathcal{K}} [\mathcal{D}^{f_k}(1^n) = 1] - \Pr_{f \xleftarrow{\$} \{\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}\}} [\mathcal{D}^f(1^n) = 1] \right| \leq \epsilon(n) \quad (3.1)$$

Consider, for example, the following **SKES** using a pseudorandom function, as found in Proposition 5.4.18 in [Gol04]. In this scheme, the pseudorandom function is used to both encrypt and decrypt messages using the same key.

Construction 3.8. Upon a security parameter n , let $\mathcal{K} = \{0, 1\}^n$ be a key space and let $\{f_k\}_{k \in \mathcal{K}}$ be a family of keyed functions over \mathcal{K} , where $f_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$. Then, let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be the symmetric-key encryption scheme defined as follows:

1. (key generation) **KeyGen**: on input 1^n , generate a key $k \xleftarrow{\$} \{0, 1\}^n$;
2. (encryption) **Enc** _{k : on message m , choose a randomness $r \xleftarrow{\$} \{0, 1\}^n$ and output a ciphertext $\text{Enc}_k(m; r) = (r, f_k(r) \oplus m)$;}
3. (decryption) **Dec** _{k : on cipher (r, c) , output $\text{Dec}_k(r, c) = c \oplus f_k(r)$;}
4. (correctness) $(\text{Dec}_k \circ \text{Enc}_k)(m; r) = (f_k(r) \oplus m) \oplus f_k(r) = m$.

In fact, this scheme already satisfies the notion of IND-CPA security, for example as shown in [KL15]. In Chapter 9, we introduce a class of quantum-secure pseudorandom functions and prove the IND-CCA1 security of this scheme, even in a setting in which the adversary is given quantum superposition access to the encryption oracle Enc_k and decryption procedure Dec_k .

In the next section, we provide a formal definition of the Learning with Errors problem, as introduced in [Reg09].

3.5 Learning with Errors

The Learning with Errors problem can be stated in multiple variants, such as the search problem or the decision problem. In the following, we begin by first defining the Learning with Errors search problem, as introduced in the introductory section.

Definition 3.9 (LWE Problem).

Let $n \geq 1$ be a security parameter, let q be a prime and let χ be a discrete probability distribution over errors in $\mathbb{Z}/q\mathbb{Z}$. Let $s \in (\mathbb{Z}/q\mathbb{Z})^n$ be a secret string and let $A_{s,\chi}$ be the probability distribution on $(\mathbb{Z}/q\mathbb{Z})^n \times \mathbb{Z}/q\mathbb{Z}$ that performs the following:

1. Sample a uniformly random string $a \in (\mathbb{Z}/q\mathbb{Z})^n$.
2. Sample an error $e \in \mathbb{Z}/q\mathbb{Z}$ according to error distribution χ_q .
3. Output $(a, \langle a, s \rangle + e)$, such that all additions are performed in $\mathbb{Z}/q\mathbb{Z}$ with respect to the modulus q .

We say that a PPT algorithm \mathcal{A} solves the Learning with Errors problem $\mathsf{LWE}_{q,\chi}$ with modulus q and error distribution χ if, for any $s \in (\mathbb{Z}/q\mathbb{Z})^n$ and an arbitrary number of independent noisy samples from $A_{s,\chi}$, \mathcal{A} outputs the secret s with nonnegligible probability.

Typically one chooses an error distribution $\chi_{\eta,q} \sim \mathcal{N}(0, \eta^2 q^2)$ that follows a discrete Gaussian distribution rounded to the nearest integer and reduced modulo q , where the noise magnitude $\eta > 0$ is taken to be $1/\text{poly}(n)$. Chebyshev's inequality allows us to conveniently control the standard deviation ηq towards a sharply peaked error distribution around the origin for an appropriate choice of parameters η and q . As Regev argues, there are several reasons that speak in favor of the hardness of the LWE problem, particularly its close relationship to lattice-problems and the *Learning Parity with Noise* problem [CSS14], both studied extensively and believed to be hard. Since LWE can be thought of as a generalization of the LPN problem, we believe that LWE must also be hard. Furthermore, the best known classical algorithms for solving the LWE problem so far run in exponential time [BKW03].

3.5.1 Decision Learning with Errors

A related variant of the LWE problem is found in the task of determining whether a given sample results from a noisy linear equation on a secret string, or a genuine uniformly random sample.

Definition 3.10 (Decision LWE).

Let $\text{LWE}_{q,\chi}$ be given by a sampling probability distribution $A_{s,\chi}$ for a string $s \in (\mathbb{Z}/q\mathbb{Z})^n$ and let U be the uniform distribution over $(\mathbb{Z}/q\mathbb{Z})^n \times \mathbb{Z}/q\mathbb{Z}$. We say that $\text{LWE}_{q,\chi}$ satisfies the decisional LWE assumption ($\text{DecLWE}_{q,\chi}$) with modulus q and error distribution χ if, for all PPT distinguishers \mathcal{D} , there exists a negligible function $\epsilon(n)$ such that:

$$\left| \Pr_{s \xleftarrow{\$} (\mathbb{Z}/q\mathbb{Z})^n} [\mathcal{D}^{A_{s,\chi}}(1^n) = 1] - \Pr[\mathcal{D}^U(1^n) = 1] \right| \leq \epsilon(n), \quad (3.2)$$

where U outputs uniform samples $(a, u) \xleftarrow{\$} (\mathbb{Z}/q\mathbb{Z})^n \times \mathbb{Z}/q\mathbb{Z}$.

Remarkably, as Oded Regev showed, there exists a simple reduction of the LWE search problem towards the decisional LWE problem. While it is clear that an efficient algorithm for the search LWE problem implies the existence of an algorithm for the decisional LWE problem, the opposite implication is guaranteed by the following lemma:

Lemma 3.11 ([Reg09], Decision LWE to Search LWE).

Let $n \geq 1$ be a security parameter, q be a prime and let $A_{s,\chi}$ be a sampling probability distribution $A_{s,\chi}$ for a string $s \in (\mathbb{Z}/q\mathbb{Z})^n$ and discrete probability distribution $\chi_{q,\eta}$ over errors in $\mathbb{Z}/q\mathbb{Z}$. If \mathcal{A} is an algorithm that solves the $\text{DecLWE}_{q,\chi}$ problem with nonnegligible probability over a uniform choice of strings s , then there exists an efficient algorithm \mathcal{A}' receiving samples from $A_{s,\chi}$ that solves the search LWE problem with probability exponentially close to 1.

3.5.2 Symmetric-Key Constructions and Security

Let us now consider the following symmetric-key encryption scheme motivated by the LWE hardness assumption, as suggested in [Reg05].

Construction 3.12 (LWE-SKES).

Consider the symmetric-key encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$, defined by:

1. (key generation) on input 1^n , KeyGen generates a key $s \xleftarrow{\$} (\mathbb{Z}/q\mathbb{Z})^n$;
2. (encryption) for each bit $b \in \{0, 1\}$, use $A_{s,\chi}$ to encrypt as follows:

$$\text{Enc}_s(b) = (a, \langle s, a \rangle + b \cdot \left\lfloor \frac{q}{2} \right\rfloor + e); \quad (3.3)$$

3. (decryption) upon cipher (a, c) , output 0 if the outcome of $\text{Dec}_s(a, c) = c - \langle a, s \rangle$ is closer to 0 than $\left\lfloor \frac{q}{2} \right\rfloor$, else output 1.
4. (correctness) $(\text{Dec}_s \circ \text{Enc}_s)(b) = b$ (with high probability as long as $e < \left\lfloor \frac{q}{4} \right\rfloor$)

Using the decisional LWE assumption, we can easily show that LWE-SKES indeed satisfies a notion of indistinguishability under a chosen-plaintext attack.

Theorem 3.13. *Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be the LWE-SKES encryption scheme from Construction 3.12. Then Π is IND-CPA-secure.*

Proof. We introduce a hybrid game by modifying the security game in a way that is indistinguishable (to any PPT adversary) from the original game in order to arrive at a security game in which the challenge is perfectly hidden.

GAME 0: In the standard hybrid, the adversary is playing the IND-CPA security game for the original scheme Π in Construction 3.12. Prior to the challenge, the adversary chooses message bits b_0, b_1 and is given access to an encryption oracle $\text{Enc}_s(\cdot)$. Upon receiving a challenge cipher $(a^*, c^*) \leftarrow \text{Enc}_s(b)$, the adversary may perform additional queries to the encryption oracle and the goal is to decide whether the challenge corresponds to an encryption of b_0 or b_1 .

GAME 1: In the this hybrid, the challenger instead responds with uniformly random samples $(a, c) \xleftarrow{\$} (\mathbb{Z}/q\mathbb{Z})^n \times \mathbb{Z}/q\mathbb{Z}$ upon each encryption query, as well as with a challenge $(a^*, c^*) \xleftarrow{\$} (\mathbb{Z}/q\mathbb{Z})^n \times \mathbb{Z}/q\mathbb{Z}$. From the decisional LWE assumption in Definition 3.10 and Lemma 3.11, it follows that no PPT adversary can distinguish between genuine LWE samples or uniformly random samples (both with $b \lfloor \frac{q}{2} \rfloor$ added to them).

Since adopting this hybrid game only negligibly affects the success probability of any PPT adversary, we arrive at a security game in which the adversary cannot win, except with at most negligible probability better than guessing at random. \square

3.5.3 Separation Result

Let us now conclude this chapter with a simple separation between the two notions of security from Section 3.3 and show that there exist schemes that are IND-CPA-secure but not IND-CCA1-secure. Using the LWE-SKES scheme, we can easily prove such a separation. The intuition is that decryption oracle access in this scheme allows the adversary to evaluate the noisy inner product upon arbitrary inputs. As a result, the adversary can determine parts of the secret key using only a few queries to its decryption oracle.

Lemma 3.14. *Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be the LWE-SKES encryption scheme from Construction 3.12. Then Π does not satisfy IND-CCA1-security.*

Proof. In the IND-CCA1 security game for Π , the adversary chooses message bits b_0, b_1 and is given access to an encryption oracle $\text{Enc}_s(\cdot)$, as well as a decryption oracle $\text{Dec}_s(\cdot)$. Upon receiving a challenge cipher $(a^*, c^*) \leftarrow \text{Enc}_s(b)$, the adversary may perform additional queries to the encryption oracle (but not to the decryption oracle) and the goal is to decide whether the challenge corresponds to an encryption of b_0 or b_1 . In order to query the decryption oracle, the adversary may choose a pair $(a, c) \in (\mathbb{Z}/q\mathbb{Z})^n \times \mathbb{Z}/q\mathbb{Z}$ and receive an output $\text{Dec}_s(a, c) = c - \langle a, s \rangle = b \lfloor \frac{q}{2} \rfloor + e$ for an unknown error e . Depending on whether the output is closer to 0 or to $\lfloor \frac{q}{2} \rfloor$, the adversary can guess the underlying encryption with

high probability. This suggests the following attack in order to determine the secret key s : The adversary queries its decryption oracle onto pairs (e_i, c) , where $e_i = (0, \dots, 1, \dots, 0)$ and where the i -th index is 1 and e_i is 0 everywhere else. Thus, $\text{Dec}_s(e_i, c) = c - \langle e_i, s \rangle = b \lfloor \frac{q}{2} \rfloor + e$. Next, the adversary computes $c - \text{Dec}_s(e_i, c) = s_i + e$. By repeating this procedure and computing $X_i = c - \text{Dec}_s(e_i, c)$ a number of times, the adversary can average out the noise for the i -th component of s , thereby obtaining a nonnegligible advantage in the security game. \square

4 Quantum Computation

Quantum information processing is concerned with the storage and manipulation of information in a quantum system. The fundamental unit of information is the *qubit*, a quantum two-level system of states $|0\rangle$ and $|1\rangle$. Fortunately, nature presents us with many ways of realizing a qubit in a physical system. Typical representations of a qubit are found in the two spin 1/2 states of a particle, the vertical or horizontal polarization of a photon or simply a ground and excited state in the energy spectrum of an atom. In this chapter, we give a brief overview of the most important concepts in the theory of quantum computing to date. For further reading, we refer to [NC10]. With regard to the physical realization of quantum computers, we refer to Chapter 10.

4.1 Formalism

A quantum system is a Hilbert space \mathcal{H} , a complex vector space together with an inner product $\langle \cdot | \cdot \rangle$. A qubit is a quantum system $|\psi\rangle$ of mutually orthogonal basis states $|0\rangle$ and $|1\rangle$, given by a normalized state vector of amplitudes $|\alpha|^2 + |\beta|^2 = 1$, where

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle. \quad (4.1)$$

Contrary to classical bits of information that carry definite states of either 0 or 1, a qubit can be represented as a continuous superposition of two basis states. By introducing angular degrees of freedom ϕ and θ , a qubit can be visualized as a point on the *Bloch sphere*, as in Figure 1, and written as²

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle. \quad (4.2)$$

Given two quantum systems \mathcal{H}_A and \mathcal{H}_B , the composition results in a joint quantum system given by $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$, the tensor product of the two systems. Thus, for $|\psi\rangle_A \in \mathcal{H}_A$ and $|\phi\rangle_B \in \mathcal{H}_B$, the product state is given by $|\psi\rangle_A \otimes |\phi\rangle_B$. For example, if $|\psi\rangle_A = \alpha|0\rangle + \beta|1\rangle$ and $|\phi\rangle_B = \delta|0\rangle + \gamma|1\rangle$, then:

$$|\psi\rangle_A \otimes |\phi\rangle_B = \alpha\delta|0\rangle \otimes |0\rangle + \alpha\gamma|0\rangle \otimes |1\rangle + \beta\delta|1\rangle \otimes |0\rangle + \beta\gamma|1\rangle \otimes |1\rangle. \quad (4.3)$$

For the sake of brevity, we often write $|\psi\rangle|\phi\rangle = |\psi\rangle_A \otimes |\phi\rangle_B$. Furthermore, we shall also frequently adopt the notation $|00\rangle$ instead of $|0\rangle|0\rangle$, as well as $|01\rangle$, $|10\rangle$ and $|11\rangle$. This allows us to conveniently represent $|\psi\rangle|\phi\rangle$ using a decimal instead of a binary expression:

$$\alpha_0|0\rangle + \alpha_1|1\rangle + \alpha_2|2\rangle + \alpha_3|3\rangle. \quad (4.4)$$

In general, a collection of n qubits forms a *register* of size n :

$$|\Psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x|x_1\rangle|x_2\rangle \dots |x_n\rangle, \quad (4.5)$$

²Note that we ignore the contributions from an overall phase as it produces no observable effects.

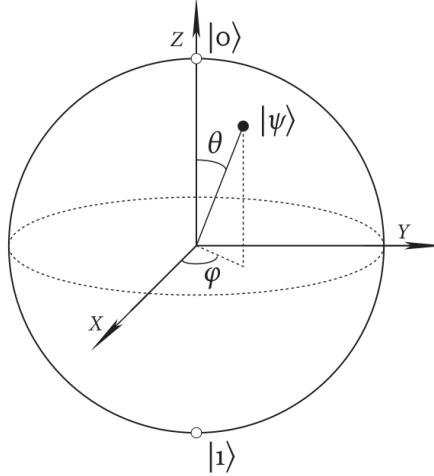


Figure 1: ([Wil13]) The Bloch sphere.

where, due to normalization, we require $\sum_x |\alpha_x|^2 = 1$. Equivalently, we can also consider the above as a superposition of 2^n different states in a decimal expression:

$$|\Psi\rangle = \sum_{x=0}^{2^n-1} \alpha_x |x\rangle. \quad (4.6)$$

Excluding the overall phase, the description of an n -qubit state already requires an enormous amount of $2^n - 1$ many complex numbers, as $\mathcal{H}_2^n \cong \mathbb{C}^{2^n}$. This fact can be exploited in quantum parallelism, which we discuss in the subsequent chapters. More generally, for $d \geq 2$, it is also useful to consider *qudits*, a quantum system of computational states $|0\rangle, |1\rangle, \dots, |d-1\rangle$ in a register of size n :

$$|\Psi\rangle = \sum_{x \in (\mathbb{Z}/d\mathbb{Z})^n} \alpha_x |x_1\rangle |x_2\rangle \dots |x_n\rangle. \quad (4.7)$$

Similarly, by adopting a decimal expression, we can write:

$$|\Psi\rangle = \sum_{x=0}^{d^n-1} \alpha_x |x\rangle. \quad (4.8)$$

In this case, the computational space $\mathcal{H}_d^n \cong \mathbb{C}^{d^n}$ features an enormous amount of $q^n - 1$ different states. The use of qudits is particularly useful in the context of the LWE problem of the later sections. While qudits are certainly more difficult to realize in a physical system, they can easily be emulated with qubits by using a block encoding in which each qudit is packed into $\lceil \log_2(d) \rceil$ many qubits.

A quantum system with a well-defined state vector $|\psi\rangle$ in \mathcal{H} is said to be *pure*. The most general state of a quantum system, however, is a *mixed* state described by a *density operator* $\rho \in \mathcal{D}(\mathcal{H})$, the set of positive semidefinite Hermitian matrices of trace equal to 1. We can interpret the density operator as a statistical ensemble of pure states $|\psi_i\rangle$, where

$\sum_i p_i = 1$, $p_i \geq 0$ and

$$\rho = \sum_i p_i |\psi_i\rangle \langle \psi_i|. \quad (4.9)$$

If ρ is pure, then ρ has rank 1 and we can conveniently write $\rho = |\psi\rangle \langle \psi|$. Furthermore, we can distinguish between pure and mixed states by using the fact that $\text{tr}(\rho^2) = 1$, if and only if ρ is pure, whereas $\text{tr}(\rho^2) < 1$, if and only if ρ is mixed.

4.2 Unitary Evolution

In the previous section, we introduced the concept of a *qubit*, a quantum system $|\psi\rangle$ described by a continuous superposition of states $|0\rangle$ and $|1\rangle$. Computation, understood as simply the manipulation of encoded information such as bits, requires a notion of what transformations are possible within a certain model of computation. Just as in Turing's abstract model of computation, it is necessary to define a model together with a set of rules on how to operate symbols stored on the equivalent of a tape by a set of instructions. In order to define what computation means in the quantum model of computation, we require one of the postulates of quantum mechanics:

The time evolution of a closed quantum system is governed by the *Schrödinger equation*,

$$i\hbar \frac{d|\psi\rangle}{dt} = \mathcal{H}|\psi\rangle, \quad (4.10)$$

where \hbar is *Planck's constant* and \mathcal{H} is the Hamiltonian operator of the system. If the Hamiltonian is time-independent, the Schrödinger equation gives rise to the following dynamics of the state vector:

$$|\psi(t)\rangle = \exp\left(\frac{-i\mathcal{H}t}{\hbar}\right) |\psi(0)\rangle. \quad (4.11)$$

The associated time-evolution operator,

$$U = \exp\left(\frac{-i\mathcal{H}t}{\hbar}\right), \quad (4.12)$$

is a *unitary* evolution operator, i.e. satisfies $UU^\dagger = \mathbb{1}$, and allows us to write (4.11) as:

$$|\psi(t)\rangle = U|\psi(0)\rangle. \quad (4.13)$$

Consequently, we can also write the unitary evolution of a density operator as

$$\rho(t) = \sum_i p_i |\psi_i(t)\rangle \langle \psi_i(t)| = \sum_i p_i U |\psi_i(0)\rangle \langle \psi_i(0)| U^\dagger = U\rho(0) U^\dagger. \quad (4.14)$$

Since an ideal qubit is required to be a closed quantum system, any unitary time-evolution describing a computation corresponds to a rotation on the Bloch sphere. Furthermore, the time-evolution of a quantum system under a given stationary Hamiltonian is reversible through its Hermitian adjoint U^\dagger . Consequently, all unitary quantum gates must be inherently reversible. As we discuss in the next sections, this fact has important consequences for many elementary operations.

4.3 Quantum Measurement

The measurement postulate of quantum mechanics specifies how information is retrieved in the quantum model of computation. Thus, in accordance with the laws of quantum mechanics, a measurement of a quantum state translates into classical measurement outcomes according to a set of rules. In this section, we highlight the most relevant notions of measurement required for the work contained in this thesis.

Quantum measurements are described by a set of *measurement operators* $\{M_m\}$ acting on the state space of a given system. These operators obey a completeness relation $\sum_m M_m^\dagger M_m = \mathbb{1}$, where m labels the measurement outcome of the associated measurement operator. Let $|\psi\rangle$ be the state vector of the system prior to measurement. Then, the probability that outcome m occurs is:

$$p(m) = \langle \psi | M_m^\dagger M_m | \psi \rangle. \quad (4.15)$$

The post-measurement state is subsequently renormalized and given by:

$$\frac{M_m |\psi\rangle}{\sqrt{\langle \psi | M_m^\dagger M_m | \psi \rangle}}. \quad (4.16)$$

For example, given the qubit from the previous sections,

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad (4.17)$$

a measurement in the *computational basis* is defined by two measurement operators, where $M_0 = |0\rangle \langle 0|$ and $M_1 = |1\rangle \langle 1|$. Each measurement operator is Hermitian, since $M_0^2 = M_0$ and $M_1^2 = M_1$, so that the completeness relation is obeyed. The probabilities of the respective outcomes are given by:

$$p(0) = \langle \psi | M_0^\dagger M_0 | \psi \rangle = \langle \psi | 0 \rangle \langle 0 | \psi \rangle = |\alpha|^2 \quad (4.18)$$

$$p(1) = \langle \psi | M_1^\dagger M_1 | \psi \rangle = \langle \psi | 1 \rangle \langle 1 | \psi \rangle = |\beta|^2. \quad (4.19)$$

Consequently, a measurement results in $|0\rangle$ with probability $|\alpha|^2$, and $|1\rangle$ with probability $|\beta|^2$. This brings us to a special case of measurements, the class of *projective measurements*. Here, the measurement operators are given by hermitian operators $\{P_m\}$, so-called *projectors*, that obey a completeness relation $\sum_m P_m = \mathbb{1}$ and satisfy $P_n P_m = \delta_{n,m} P_m$.

The probability to observe the outcome m is given by:

$$p(m) = \langle \psi | P_m | \psi \rangle, \quad (4.20)$$

whereas the post-measurement state is

$$\frac{P_m |\psi\rangle}{\sqrt{\langle \psi | P_m | \psi \rangle}}. \quad (4.21)$$

A final class of more general measurements we consider is that of *POVM measurements* (Positive-Operator-Valued Measure) [NC10], where the post-measurement state is of little interest and the concern lies on the outcome probabilities corresponding to a set of measurement operators. In this context, a set of positive semidefinite measurement operators $\{E_m\}$ is employed such that $\sum_m E_m = \mathbb{1}$.

4.4 Universal Quantum Gates

In this section we introduce elementary quantum gates, in particular those that allow for universal quantum computation. In Section 4.2, we observed that all quantum gates must correspond to unitary transformations, and are thus inherently reversible. While classical universality of logic gates is achieved by using only a NAND gate, we require a certain universal set of at least three elementary gates for quantum computation.

Let us begin with a few examples of single-qubit quantum gates. A simple set of single-qubit gates are the X, Y, Z -gates, resembling the *Pauli matrices* σ_x, σ_y and σ_z :

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (4.22)$$

Consider, for example, a qubit $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. In vector representation, we compute the action of the X -gate as follows:

$$X|\psi\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \cdot \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = \beta|0\rangle + \alpha|1\rangle. \quad (4.23)$$

One of the most frequently used operations in quantum computing is that of the *Hadamard gate*, which is given by

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \quad (4.24)$$

The Hadamard gate, often described as a square root of the X -gate, completes only half of a 180° rotation on the Bloch sphere and maps the basis states onto an equal superposition and back:

$$|0\rangle \xleftrightarrow{H} |+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \quad |1\rangle \xleftrightarrow{H} |-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}} \quad (4.25)$$

Another important single-qubit gate is the *phase-shift gate* in which ϕ denotes the angle of rotation. The special case where $\phi = \pi/4$ is often referred to as the T-gate:

$$\phi = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\phi} \end{bmatrix}, \quad T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}. \quad (4.26)$$

In addition, we consider the rotation operators around the x, y and z axis:

$$R_x(\theta) = \begin{bmatrix} \cos \frac{\theta}{2} & -i \sin \frac{\theta}{2} \\ -i \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{bmatrix} \quad R_y(\theta) = \begin{bmatrix} \cos \frac{\theta}{2} & -\sin \frac{\theta}{2} \\ \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{bmatrix} \quad R_z(\theta) = \begin{bmatrix} e^{-i\frac{\theta}{2}} & 0 \\ 0 & e^{i\frac{\theta}{2}} \end{bmatrix}. \quad (4.27)$$

In fact, *any* unitary single-qubit operation U can be decomposed using the rotation operators above.

Theorem 4.1 ([NC10], Theorem 4.10).

The two rotation operations R_x and R_y comprise a basis for all single-qubit operations: For every 2×2 unitary operation U , there exist real numbers α, β, γ and δ such that:

$$U = e^{i\alpha} R_x(\beta) R_y(\gamma) R_x(\delta).$$

Let us now conclude our discussion on quantum gates with two-qubit gates, perhaps the most striking class of operations found in quantum computers. Early work by Deutsch, Eckert and Josza suggests that this class of gates is precisely the set of operations that entangles qubits with one another, thereby providing the foundation for most quantum computations. The most important two-qubit gate is the *controlled-NOT* (CNOT) gate, an operation that performs a bit flip on a target bit if and only if the control qubit is $|0\rangle$. Another important gate is the Toffoli (CCNOT) gate, a three-qubit gate that flips the last qubit only if and only if all three inputs correspond to $|1\rangle$. The matrix representations are given by:

$$\text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad \text{Toffoli} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (4.28)$$

Equivalently, these two-qubit and three-qubit gates can also be written as the following operations:

$$\text{CNOT: } |x\rangle|y\rangle \longrightarrow |x\rangle|x\oplus y\rangle \quad (4.29)$$

$$\text{Toffoli: } |x\rangle|y\rangle|z\rangle \longrightarrow |x\rangle|y\rangle|z\oplus x\wedge y\rangle, \quad (4.30)$$

where \oplus denotes addition modulo 2 and \wedge denotes the AND operation (Table 1). Finally, we also consider the *controlled-Z* (CZ) gate, an operation that features an additional minus sign and has the following matrix representation:

$$\text{CZ} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \quad (4.31)$$

The following theorem states ensures that quantum computation is indeed universal using only a limited set of gates.

Theorem 4.2 (Universal set of quantum operations, [Deu89] [Kit97]).

The Hadamard gate, the Toffoli gate and phase-shift gate comprise a universal basis for any quantum operation: For every $D \geq 3$, there exists $l \leq 100(D \log \frac{1}{\epsilon})^3$ such that every unitary $D \times D$ matrix U can be approximated by a sequence of the above unitary gates to $\epsilon > 0$ degree of accuracy:

$$|U_{i,j} - (U_l \cdots U_1)_{i,j}| < \epsilon,$$

where the index (i, j) denotes the entries of the respective matrices.

While a basis for quantum computation is not limited to precisely the set as given in Theorem 4.2, it nevertheless provides a convenient choice of elementary gates.

4.5 The No-Cloning Theorem

One of the most defining aspects of quantum information is the fact that it cannot be copied. The *No-Cloning Theorem* is attributed to Wootters and Zurek [WZ82] and can in brief be stated as follows:

There exists no universal quantum operation that can make an identical copy of an unknown quantum state.

We give a short proof by contradiction, as for example found in [Wil13], that illustrates the impossibility of such an operation: Suppose there exists a two-qubit operator U with the above properties. Thus, if we take an arbitrary qubit $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ as input, we expect:

$$U|\psi\rangle|0\rangle \rightarrow |\psi\rangle|\psi\rangle \quad (4.32)$$

$$= (\alpha|0\rangle + \beta|1\rangle)(\alpha|0\rangle + \beta|1\rangle) \quad (4.33)$$

$$= \alpha^2|0\rangle|0\rangle + \alpha\beta|0\rangle|1\rangle + \beta\alpha|1\rangle|0\rangle + \beta^2|1\rangle|1\rangle. \quad (4.34)$$

Due to the universality of the U operation, we also have that:

$$U|0\rangle|0\rangle \rightarrow |0\rangle|0\rangle \quad (4.35)$$

$$U|1\rangle|0\rangle \rightarrow |1\rangle|1\rangle. \quad (4.36)$$

However, we must now conclude that U violates the linearity of quantum mechanics, since

$$U(\alpha|0\rangle + \beta|1\rangle)|0\rangle = \alpha^2|0\rangle|0\rangle + \beta^2|1\rangle|1\rangle, \quad (4.37)$$

where, for general α and β , these two expressions (4.32) and (4.37) are not equal.

4.6 The Quantum Circuit Model

A quantum computation typically starts out in some initial state $|00\dots0\rangle$, then performs a sequence of single-qubit and two-qubit gates and finally ends with a measurement in the computational basis. The *quantum circuit model* provides us with a convenient way of representing any quantum computation pictorially.

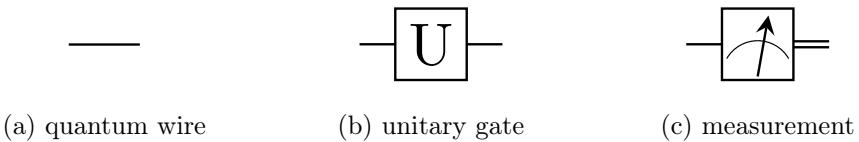


Figure 2: The quantum circuit model. Quantum wires (a) represent the history of single qubits in time progressing from left to right. Single-qubit unitary gates (b) are represented as a box applied to a single quantum wire. Measurements (c) of qubits are provided by projective measurements in the computational basis.

In the following, we consider the representations of a Hadamard gate acting on an initial state $|0\rangle$, as well as a CNOT gate and a CZ gate.

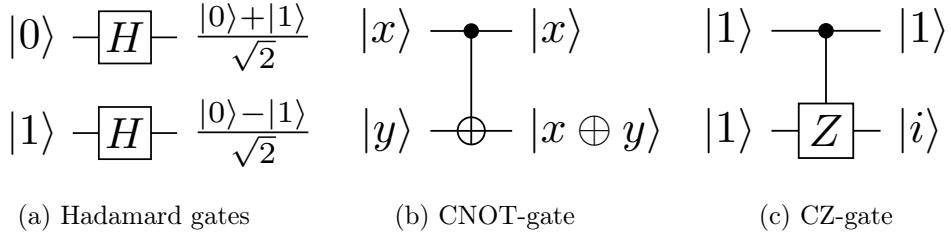


Figure 3: Single-qubit gates and two-qubit gates. The Hadamard gates (a) each act on a single quantum wire. The CNOT gate (b) adds the value of the control qubit into the target qubit. The CZ-gate (c) performs a phase flip only if both the control and target qubit are $|1\rangle$.

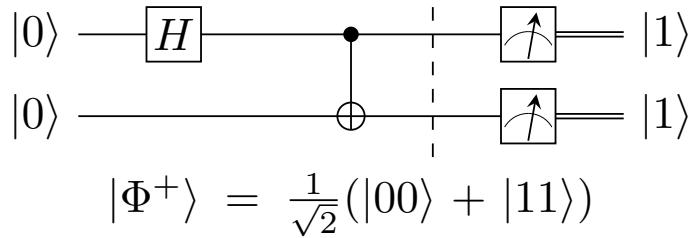


Figure 4: A quantum circuit that prepares an entangled state $|\Phi^+\rangle$ (as indicated by the dashed line), the famous Einstein-Podolsky-Rosen (EPR)-pair $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. Here, a final measurement in the computational basis results in the outcome $|1\rangle |1\rangle$.

4.7 Quantum Parallelism

In the previous sections, we recognized that unitary quantum gates are inherently reversible. Can we, nevertheless, still simulate classical computation using only reversible gates? Consider for example, the classical NAND gate, as shown in Table 1. The NAND-gate, a universal logic gate for classical computation, is inherently irreversible. Knowing that the output is 1, we cannot conclude with certainty whether the input was in fact 00, 01 or 10. More generally, consider the problem of computing the following transformation:

$$x \xrightarrow{f} f(x) \quad (4.38)$$

If f is a bijective operation, we can reverse this transformation and recover the input. However, if f is irreversible, we can attach the input and still achieve an overall reversible operation, as follows:

$$(x, y) \xrightarrow{f} (x, y \oplus f(x)). \quad (4.39)$$

Note that, when performing this operation twice, we obtain the original input pair. A well known trick due to Charles Bennet allows us to compute irreversible transformations using only reversible quantum gates at the expense of a few additional registers. By simply attaching additional input registers prior to the evaluation of the function, a reversible unitary transformation U_f is possible in which later registers can be uncomputed. As a

Inputs:		Outputs:	
A	B	A AND B	A NAND B
0	0	0	1
1	0	0	1
0	1	0	1
1	1	1	0

Table 1: Two classical logic gates: AND(\wedge) and NAND(\uparrow).

result, this allows us to define operations, such as:

$$|x\rangle|y\rangle \xrightarrow{U_f} |x\rangle|y \oplus f(x)\rangle, \quad (4.40)$$

where $U_f^\dagger U_f = \mathbb{1}$. This subsumes an evaluation of f , as we can initialize the second qubit to $y = 0$ and compute the output of f as follows:

$$|x\rangle|0\rangle \xrightarrow{U_f} |x\rangle|f(x)\rangle \quad (4.41)$$

One of the essential features of quantum algorithms is quantum parallelism, the ability to prepare a superposition of states for simultaneous evaluation.

Consider a simple boolean function $f : \{0, 1\} \rightarrow \{0, 1\}$ and suppose we have access to a unitary gate that evaluates f onto the presented inputs, such as in (4.40). By preparing two initial states $|0\rangle|0\rangle$ and applying a single Hadamard gate onto the first register, we can evaluate f and exploit quantum parallelism (Figure 5). The result of such a transformation

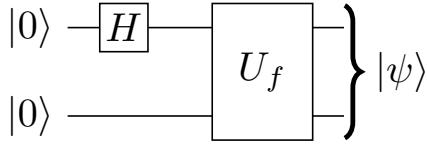


Figure 5: A quantum circuit that prepares a superposition $|\psi\rangle$ which simultaneously evaluates f on both 0 and 1.

is an output state:

$$|\psi\rangle = \frac{1}{\sqrt{2}} \sum_{x \in \{0,1\}} |x\rangle|f(x)\rangle = \frac{1}{\sqrt{2}}(|0\rangle|f(0)\rangle + |1\rangle|f(1)\rangle). \quad (4.42)$$

Remarkably, the transformation achieves a superposition that contains information on both $f(0)$ and $f(1)$. Simultaneous evaluation of a function is what gives power to quantum parallelism: a single quantum evaluation of a function can result in a state that features evaluations of f in superposition. However, if one were to measure the state in the computational basis, it would simply collapse and reveal only a random evaluation of the function. Notice that, by generalizing this to a collection of n -qubits, the amount of evaluations of the function grows exponentially in n . Consequently, research on quantum algorithms concerns techniques that exploit such information hidden in the superposition to one's advantage.

4.8 Decoherence

As with any physical implementation of a computing device, not all operations can be done perfectly and there remains an unavoidable risk of error. Typically, one distinguishes between two classes of errors. We encounter both *memory errors* that occur during storage of information, as well as *operational errors* that occur during manipulation of stored information. In the section on error correcting codes, we provide further discussion on how to correct for such errors.

4.8.1 Quantum Noise Models

In order to develop a successful noise model, we follow [NC10] and adopt a theory of quantum channels and the operator-sum-representation. In this framework, we consider noise models as *discrete state changes* without reference to time.

For the remainder of the section, we let ρ be a quantum system of computational states $|0\rangle$ and $|1\rangle$ and define the action of a noisy quantum channel by a completely positive and trace preserving (CPTP) operation \mathcal{E} , where for operation elements E_0 and E_1 , we have:

$$\rho \longrightarrow \mathcal{E}(\rho) = E_0 \rho E_0^\dagger + E_1 \rho E_1^\dagger. \quad (4.43)$$

A simple quantum noise channel is the bit-flip channel that, with probability $\eta > 0$, maps the state $|0\rangle$ to the state $|1\rangle$ and vice-versa:

$$E_0 = \sqrt{1-\eta} \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad E_1 = \sqrt{\eta} \cdot \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}. \quad (4.44)$$

Thus, in the operator-sum-representation, we can now write the bit-flip channel as:

$$\rho \longrightarrow \mathcal{E}_X(\rho) = (1-\eta)\rho + \eta X \rho X^\dagger, \quad (4.45)$$

where X corresponds to the bit-flip gate from the earlier sections.

Consider now the case of a single qubit, a quantum system that starts out in a pure state $\rho = |\psi\rangle\langle\psi|$, where $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. Under the bit-flip channel \mathcal{E} , the state evolves into a statistical mixture according to (4.45). Therefore, with probability $1-\eta$, we recover the original state,

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad (4.46)$$

and, with probability η , we find the system in a state:

$$|\psi^\perp\rangle = \beta|0\rangle + \alpha|1\rangle. \quad (4.47)$$

This noise model is often called *classification noise* and will be highly relevant for the quantum learning algorithms of the later sections.

Another elementary quantum noise channel is the phase-flip channel that, with probability $\eta > 0$, maps the state $|1\rangle$ to the state $-|1\rangle$, where

$$E_0 = \sqrt{1-\eta} \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad E_1 = \sqrt{\eta} \cdot \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (4.48)$$

Thus, in the operator-sum-representation, we can also write the phase-flip channel as:

$$\rho \longrightarrow \mathcal{E}_Z(\rho) = (1 - \eta)\rho + \eta Z\rho Z^\dagger, \quad (4.49)$$

where Z corresponds to the phase-flip gate from the earlier sections.

A much more ideal noise model for quantum computation is the amplitude damping channel. This noise channel \mathcal{E}_{AD} corresponds to a scenario in which a photon is spontaneously emitted with some probability γ .

$$\rho \longrightarrow \mathcal{E}_{AD}(\rho) = E_0\rho E_0^\dagger + E_1\rho E_1^\dagger, \quad (4.50)$$

where the transition matrix operators are given by:

$$E_0 = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{bmatrix}, \quad E_1 = \begin{bmatrix} 1 & \sqrt{\gamma} \\ 0 & 0 \end{bmatrix}. \quad (4.51)$$

Thus, E_1 corresponds to the emission of a photon and a quantum of energy is lost to the environment. The operator E_0 , corresponds to the case where the state remains unchanged and a photon is not yet lost but the amplitudes are adjusted appropriately.

Finally, we consider the *depolarizing channel*, a devastating type of noise model in which all quantum information is lost to the environment and the quantum state gets replaced by a maximally mixed state with some probability η :

$$\rho \longrightarrow \mathcal{E}_D(\rho) = (1 - \eta)\rho + \eta \frac{\mathbb{1}}{2}, \quad (4.52)$$

where $\frac{\mathbb{1}}{2}$ is the maximally mixed state.

4.8.2 Independent Noise Models

The most commonly adopted noise model in computational learning theory is that of independent noise. For example, in the LWE problem each state or sample is independently corrupted by some probability. Most generally in the context of qubits, we consider independent noise by extending the noise models from this section onto quantum registers.

Consider a register of n qubits $|\Psi\rangle$, where

$$|\Psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x_1\rangle |x_2\rangle \dots |x_n\rangle. \quad (4.53)$$

By, for example, extending the bit-flip channel from this section independently onto the entire register, the result is a state

$$|\Psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x_1 \oplus e_1\rangle |x_2 \oplus e_2\rangle \dots |x_n \oplus e_n\rangle, \quad (4.54)$$

where each error e_i is sampled from a Bernoulli distribution of noise rate $\eta > 0$:

$$Bern(\eta) = \begin{cases} 1, & \text{with probability } \eta \\ 0, & \text{with probability } 1 - \eta. \end{cases} \quad (4.55)$$

Note that, upon a choice of error distribution, an independent noise model also translates naturally in the context of qudits instead of qubits.

4.9 Error Correcting Codes

In his seminal 1948 paper *A Mathematical Theory of Communication* [Sha48], Claude Shannon put forward a revolutionary view on the concept information and errors in communication. Instead of investing tedious effort to avoid them on a technical level, it is not only possible, but oftentimes even favorable, to simply correct them. In the 1950s, John von Neumann developed very successful error correcting codes in order to address noise originating in the relay architecture of present computer technology. Today's transistors achieve near-perfect fault tolerance, hence error correcting codes often do not even have to be applied. By adding additional redundant information to each bit of information, one can realize error correction.

Suppose the task is to store a single bit of information for some desired time interval $T \geq 0$. Regardless of whether an operation takes place, we consider *memory errors* that can occur spontaneously. We denote the probability that such an error occurs after time T by p . A simple way to protect the storage of information against these affects is by adding redundancy. For example, consider the following three copies of each instance of the bit:

$$0 \rightarrow 000 \tag{4.56}$$

$$1 \rightarrow 111. \tag{4.57}$$

The probability that there are no errors is given by $(1 - p)^3$. Hence, after time T , the three bits 000 remain the same. The probability that there is an error in one of the bits is $3p(1 - p)^2$, resulting in either 001, 010 or 100. Finally, the probability that there are two or more errors is $3p^2(1 - p) + p^3$. The error correction scheme now works as follows: By measuring all the bits and taking the majority vote, we can easily rule out single bit errors. Our correction method thus assigns the measurement outcomes of the bitstrings as follows:

$$\{000, 001, 010, 100\} \rightarrow 0 \tag{4.58}$$

$$\{111, 110, 101, 011\} \rightarrow 1. \tag{4.59}$$

The error correction method above is correct with probability $p_c = 1 - 3p^2 + 2p^3$. Compared to the previous error probability of p , we gain as long as $p_c \geq 1 - p$. This translates into an error probability of at most $p < \frac{1}{2}$. By dividing the time interval into slices $\Delta t = T/N$ and applying the error correcting code repeatedly after each slice, one can reach arbitrarily close success probabilities which grow in N [NC10][CZ01].

Consider now the problem of quantum error correction by analogy to classical error correction. Due to the quantum nature of information, a new framework is needed to correct for errors. This becomes apparent as we consider a number of differences as compared to classical error correction. According to the No-Cloning Theorem 4.5, there is no machine that can make a copy of an unknown quantum state. Therefore, the naive attempt of simply copying quantum information in order to achieve redundancy is not possible. Moreover, a significant feature of quantum states is that the parameters describing them are continuous. Consequently, quantum noise is also continuous and requires correction to reach up to infinite precision, hence demanding unbounded resources. And finally, classical error correction

requires read-out or state detection of the bit sequence in order to detect errors and correct them. In particular, measuring a quantum state generally destroys the state and makes recovery impossible. The *quantum fault-tolerant threshold theorem* [Pre98][AB08] states that, as long as the noise level is below a certain threshold (typically around $10^{-4} - 10^{-2}$), any quantum computation can be performed with arbitrarily small error by adopting error correction. Most notably, *Shor's code* [Sho95] achieves error correction for arbitrary single-qubit errors, including bit-flips and phase errors. In the following, we give a simple example of a three qubit quantum error correcting code, known as the bit-flip code.

Suppose the task is to store a single qubit $|\psi\rangle = c_0|0\rangle + c_1|1\rangle$ in a state that is unknown to us. Analogous to the classical case, we assume that after some time T a bit flip occurs with probability p such that the state $|\psi\rangle$ is taken to the corrupted state $X|\psi\rangle$. The bit flip error thus results in a new state $|\psi^\perp\rangle = c_0|1\rangle + c_1|0\rangle$. First, we begin by preparing the following encoding into a sequence of three logical qubits

$$|0\rangle \longrightarrow |0\rangle_L \equiv |000\rangle \quad (4.60)$$

$$|1\rangle \longrightarrow |1\rangle_L \equiv |111\rangle. \quad (4.61)$$

We can realize this encoding in a quantum state $|\psi\rangle_L = c_0|000\rangle + c_1|111\rangle$. A circuit that performs this operation upon $|\psi\rangle$ is given by:

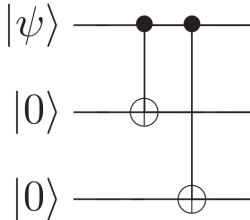


Figure 6: A quantum circuit that prepares the state $c_0|000\rangle + c_1|111\rangle$.

The probability that there are no errors in the state $|\psi\rangle_L$ is given by $(1-p)^3$. Hence, after time T , the three qubits remain the same. The probability that there is an error in just one of the qubits is $3p(1-p)^2$, resulting in either $(X \otimes \mathbb{1} \otimes \mathbb{1})|\psi\rangle_L$, $(\mathbb{1} \otimes X \otimes \mathbb{1})|\psi\rangle_L$ or $(\mathbb{1} \otimes \mathbb{1} \otimes X)|\psi\rangle_L$. Finally, the probability that there are two or more errors is $3p^2(1-p) + p^3$. The quantum error correcting code now works in two steps, as in the classical code from the previous section. First, errors are being detected and then subsequently being corrected for in the second step using a recovery procedure. To this end, consider the following sets of (incomplete) projection operators:

$$P_0 = |000\rangle\langle 000| + |111\rangle\langle 111| \quad (4.62)$$

$$P_1 = |100\rangle\langle 100| + |011\rangle\langle 011| \quad (4.63)$$

$$P_2 = |010\rangle\langle 010| + |101\rangle\langle 101| \quad (4.64)$$

$$P_3 = |001\rangle\langle 001| + |110\rangle\langle 110|. \quad (4.65)$$

The first projection operator corresponds to the case where there is no error, and the other operators correspond to a single bit-flip on one of the qubits, respectively. In any

error correcting code, the goal is to infer information on what error has occurred without destroying the superposition $|\psi\rangle$ altogether. The design of the code should therefore aim at projecting $|\psi\rangle_L$ into mutually orthogonal spaces in which we can detect the type of the error and reversibly restore the original state without at any point destroying the information. We begin by first measuring the operators above for the state $|\psi\rangle_L$. Starting with P_0 , whenever we obtain 1, we leave the state as it is knowing no error occurred. If we obtain 0, we continue by measuring the next projector P_1 . If we obtain 1, we know that a bit flip occurred on the first qubit, and we correct by applying the $(X \otimes \mathbb{1} \otimes \mathbb{1})$ operation. Similarly, we can correct for all single bit-flip errors as in the classical error correcting code and gain an advantage as long as $p < \frac{1}{2}$.

4.10 Quantum Oracles

An important abstraction in computational complexity theory and the study of decision problems is the use of oracle machines, or *oracles*. First introduced in the context of Turing machines, oracles act as a black box that assist a Turing machine in a given computational task. When presented an input x , such as an integer or a string, the oracle solves an instance of a decision problem at unit cost and returns the output $\mathcal{O}(x)$ (typically a YES/NO answer) back to the Turing machine. While this computational setting is certainly highly abstract, it does however contribute enormously to our understanding of complexity classes and is commonplace in theoretical cryptography, particularly when providing arguments for security. Oracles also turn out to be highly useful in quantum computing, especially in the study of quantum algorithms. Most notably, *Grover's search algorithm* [Gro96] relies on the use of a quantum oracle that recognizes solutions to a given search problem. Many of the earliest quantum algorithms, such as the Deutsch-Josza algorithm [DJ92], the Bernstein-Vazirani algorithm [BV93] or Simon's algorithm [Sim97] are also devised in an oracle model. In order to make an oracle evaluation reversible, we rely on the technique from the previous section. For our purposes, a quantum oracle \mathcal{O} , is a unitary operation,

$$|x\rangle |y\rangle \longrightarrow |x\rangle |y \oplus \mathcal{O}(x)\rangle, \quad (4.66)$$

acting as a black box that can be accessed by a quantum computation. The inner workings of the quantum oracle are unknown to the computation, but can evaluate upon inputs in a reversible manner, see (Figure 7). Consider now a collection of n -qubits, a quantum register

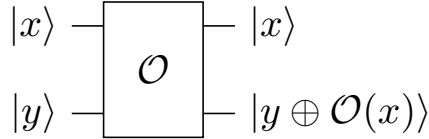


Figure 7: Quantum oracle.

of size n . As a unitary gate, the oracle obeys the linearity of quantum mechanics and can

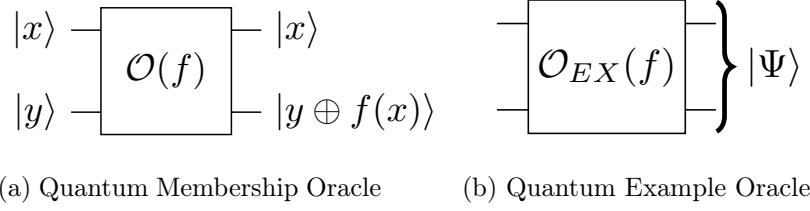


Figure 8: The quantum membership oracle (a) can be queried on arbitrary inputs during a quantum computation. The quantum example oracle (b) responds to queries by outputting a uniform superposition $|\Psi\rangle$ over all evaluations of the function.

therefore be queried on a superposition over all inputs:

$$\sum_{x,y \in \{0,1\}^n} \alpha_{x,y} |x\rangle |y\rangle \longrightarrow \sum_{x,y \in \{0,1\}^n} \alpha_{x,y} |x\rangle |y \oplus \mathcal{O}(x)\rangle \quad (4.67)$$

More generally, we also further differentiate between two variants of oracles, so-called *membership oracles* and *example oracles*.

4.10.1 Membership Oracles

Consider, for example, a boolean function given by $f : \{0,1\}^n \longrightarrow \{0,1\}$. In a membership oracle model, the oracle provides direct unitary input access to the function f which is to be evaluated. Upon input $x \in \{0,1\}^n$ and additional register $y \in \{0,1\}$, we define a membership oracle for the function f as an operation:

$$\mathcal{O}_f : |x\rangle |y\rangle \longrightarrow |x\rangle |y \oplus f(x)\rangle. \quad (4.68)$$

The oracle can thus be queried as a quantum gate at any step of a quantum computation.

4.10.2 Example Oracles

In the setting of computational learning theory from the subsequent chapters, we consider uniform example oracles $\mathcal{O}_{EX}(f)$ as a black box that only outputs uniform samples for a given function. Thus, upon each query, the oracle replies with a state:

$$|\Psi\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle |f(x)\rangle \quad (4.69)$$

In the context of quantum example oracles, we also consider samples that are corrupted by noise. In particular, as the example oracle consists of a quantum circuit that evaluates f on a superposition of all inputs, this operation is naturally prone to errors. We consider two models of noise in this setting.

First, we consider uniform example states $|\Psi\rangle$ that suffer from a bit-flip error in the final *result register* through a noise channel \mathcal{E}_X of magnitude $\eta > 0$, see (4.45). Equivalently, we can also express the fact that $|\Psi\rangle$ is turning into a mixture by again sampling an error from a Bernoulli distribution of noise parameter η , resulting in a state

$$|\Psi\rangle = \frac{1}{\sqrt{2^n}} \sum_{x,y \in \{0,1\}^n} |x\rangle |f(x) \oplus e\rangle. \quad (4.70)$$

Moreover, in a model resembling the LWE problem, we consider independent noise in the result register for each element in the superposition:

$$|\Psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x_1\rangle |x_2\rangle \dots |x_n\rangle |f(x) \oplus e_x\rangle. \quad (4.71)$$

Upon a choice of error distribution, such an independent noise model also translates naturally in the context of qudits instead of qubits.

5 Quantum Algorithms

Ever since the dawn of quantum computation, it was speculated that quantum computers could solve certain computational problems faster than any conventional classical computer. Historically, the first abstract model of universal computation was proposed by Alan Turing in his seminal 1936 paper, a discovery that henceforth greatly shaped the field of theoretical computer science. The *Church-Turing thesis* famously suggests that any model of computation appears at most as powerful as a Turing machine:

Any intuitively computable algorithmic process can be simulated efficiently by a Turing machine.

The importance of efficient algorithms is made more precise in *computational complexity* and concerns only computations of polynomial amounts of elementary operations, thus highlighting the set of problems that can be solved with a *feasible* use of computational resources. On the contrary, inefficient algorithms require superpolynomial amounts of resources (typically exponential) and become computationally infeasible as the size of the problem increases. The observation that the laws of physics are fundamentally quantum mechanical ultimately led David Deutsch to speculate on the prospect of computing devices that behaved inherently quantum mechanical. Deutsch's insights into universal quantum computation [Deu85] and the discovery of the first quantum algorithm outperforming classical counterparts [DJ92] provided unforeseen challenges for the Church-Turing principle. Shor's factoring algorithm [Sho94] provided further evidence that quantum computers could indeed solve computational problems for which no efficient classical algorithm is known. Still today, it is not clear whether a quantum model of computation is indeed capable of efficiently simulating any physical system in nature, i.e. whether a quantum extension of the original thesis, the *Quantum Church-Turing principle*, holds.

In this section, we review some of early quantum algorithms that offer substantial quantum speed-ups, such as the Deutsch-Josza algorithm [DJ92], the Bernstein-Vazirani algorithm [BV93] and Simon's algorithm [Sim97], each providing the basis for the algorithms of the later sections. In this thesis, we regard a quantum algorithm as a sequence of unitary operations, i.e. *computations*, operating on a product state space, for example $\mathcal{H} = \mathcal{H}_{\text{input}} \otimes \mathcal{H}_{\text{work}} \otimes \mathcal{H}_{\text{output}}$, in analogy to the tape of a Turing machine. Upon an input state $|\psi_0\rangle \in \mathcal{H}$, the quantum polynomial time (QPT) algorithm runs an efficient quantum circuit consisting of a sequence of unitary operations:

$$|\psi\rangle = U_T U_{T-1} U_{T-2} \dots U_1 |\psi_0\rangle, \quad (5.1)$$

where T denotes a polynomial (in terms of the dimension of \mathcal{H}) amount of operations. Hence, the algorithm generates an output state $|\psi\rangle$, typically followed by a measurement in the computational basis. In this chapter, we consider problems in a membership oracle model, as discussed in Section 4.10. Here, the goal of the algorithm is determine a secret property of a given function f by making queries to an oracle for f . Moreover, the oracle serves as a black box and can be accessed with queries $\mathcal{O}_f : |x\rangle |y\rangle \rightarrow |x\rangle |y \oplus f(x)\rangle$ at unit cost during the computations as follows:

$$|\psi\rangle = U_T \mathcal{O}_f U_{T-1} \mathcal{O}_f U_{T-2} \dots U_2 \mathcal{O}_f U_1 |\Psi_0\rangle. \quad (5.2)$$

5.1 Deutsch-Josza Algorithm

Consider the problem of determining whether a boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is either constant or balanced, i.e. 0 for half of the inputs and 1 else, as appeared in [DJ92]:

Algorithm 1 Deutsch-Josza Algorithm

Input: A quantum black box oracle \mathcal{O}_f for a boolean function f which is either constant or balanced.

Output: Outcome $|0^n\rangle$ if and only if f is constant, else f is balanced.

Procedure:

1. Initialize $(n + 1)$ -qubits $|0^n\rangle |1\rangle$ and apply the Hadamard transform $H^{\oplus(n+1)}$:

$$\rightarrow \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle |-\rangle$$

2. Query the quantum oracle, resulting in a *phase-kickback*:

$$\rightarrow \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle |-\rangle$$

3. Throw away the last register and then apply another Hadamard gate $H^{\otimes n}$:

$$\rightarrow \frac{1}{\sqrt{2^n}} \sum_{x,y \in \{0,1\}^n} (-1)^{\langle x,y \rangle} (-1)^{f(x)} |y\rangle$$

4. Measure the entire output state.

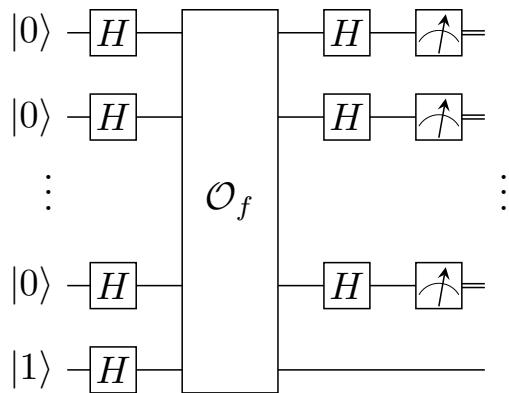


Figure 9: A quantum circuit whose outcome determines whether a boolean function is constant or balanced using only a single query to the membership oracle.

We can verify the correctness of the algorithm as follows: If we measure the final output state of the algorithm in the computational basis for the outcome $|0^n\rangle$, we observe that

$$p(0^n) = \left\| \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |0^n\rangle \right\|^2 = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} (-1)^{f(x)}. \quad (5.3)$$

Thus, the amplitudes interfere constructively towards a probability of 1 if f is constant, whereas they interfere destructively around a probability of 0 whenever f is balanced. Note that the quantum algorithm only required as much as a single query to the oracle. Classically, in the worst-case setting, any algorithm requires $\Omega(2^{n-1} + 1)$ classical queries to the oracle in order to learn more than half of the evaluations of the function.

5.2 Bernstein-Vazirani Algorithm

Another potentially interesting problem in complexity theory is the task of determining a hidden string from inner product of bit strings. In 1993, Bernstein and Vazirani [BV93] initiated the field of quantum complexity theory and proposed a quantum algorithm achieving a superpolynomial speed-up over classical algorithms. In the following, we state a popular variant of the original algorithm and discuss its speed-up. In brief, the problem can be stated as follows:

Bernstein-Vazirani Problem:

Recover a string $s \in \{0,1\}^n$ by querying an oracle for a boolean function $f_s : \{0,1\}^n \rightarrow \{0,1\}$ given by

$$f_s(x) = s_1 \cdot x_1 \oplus \dots \oplus s_n \cdot x_n = \langle s, x \rangle \bmod 2.$$

In the classical membership oracle setting, we observe that a single query to the function can only ever reveal as much as one bit of information about the secret string s . In fact, this can easily be done by considering queries on strings $e_i = (0, \dots, 1, \dots, 0)$, where the i -th index is 1 and e_i is 0 everywhere else. An algorithm performing such queries achieves an overall query complexity of $O(n)$ when determining the secret, as each iteration reveals only a single bit of the hidden string by querying

$$f_s(e_i) = \langle s, e_i \rangle \bmod 2 = s_i, \quad (5.4)$$

so that s is fully determined after a total of n queries to the function.

In the quantum membership oracle model, Bernstein and Vazirani showed that only a single quantum query to the oracle is sufficient [BV93]:

Algorithm 2 Bernstein-Vazirani Algorithm

Input: A quantum black box oracle \mathcal{O}_{f_s} for a boolean function f_s , where $f_s(x) = \langle s, x \rangle$.
The task is to determine $s \in \{0, 1\}^n$.

Output: The secret string with only a single query to the oracle.

Procedure:

1. Initialize $(n + 1)$ -qubits to $|0^n\rangle |1\rangle$ and apply the Hadamard transform $H^{\oplus(n+1)}$:

$$\rightarrow \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle |-\rangle$$

2. Query the quantum oracle, resulting in a *phase-kickback*:

$$\rightarrow \frac{1}{\sqrt{2^n}} \sum_{x,y \in \{0,1\}^n} (-1)^{\langle s, x \rangle} |x\rangle |-\rangle$$

3. Throw away the last register and then apply a Hadamard gate $H^{\otimes n}$:

$$\rightarrow \frac{1}{\sqrt{2^n}} \sum_{x,y \in \{0,1\}^n} (-1)^{\langle s, x \rangle} (-1)^{x \cdot y} |y\rangle$$

4. Measure the entire output state.

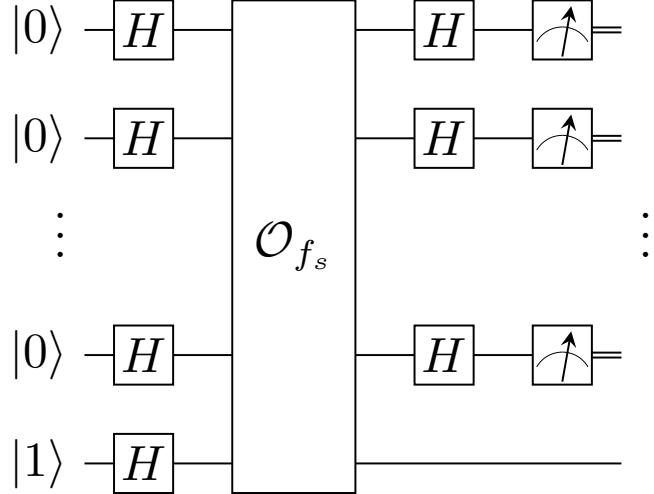


Figure 10: A quantum circuit for the Bernstein-Vazirani problem. The secret string is determined after only a single query to the membership oracle.

Thus, if we now measure the final output state of the algorithm in the computational basis for a particular outcome $m \in \{0, 1\}^n$, we observe that

$$p(m) = \left\| \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} (-1)^{\langle s, x \rangle} (-1)^{x \cdot m} |m\rangle \right\|^2 \quad (5.5)$$

$$= \frac{1}{2^n} \sum_{x \in \{0,1\}^n} (-1)^{(s_1 \oplus m_1)x_1} (-1)^{(s_2 \oplus m_2)x_2} \dots (-1)^{(s_n \oplus m_n)x_n} \quad (5.6)$$

$$= \frac{1}{2^n} \left(\sum_{x_1 \in \{0,1\}} (-1)^{(s_1 \oplus m_1)x_1} \right) \left(\sum_{x_2 \in \{0,1\}} (-1)^{(s_2 \oplus m_2)x_2} \right) \dots \left(\sum_{x_n \in \{0,1\}} (-1)^{(s_n \oplus m_n)x_n} \right) \quad (5.7)$$

$$= \frac{1}{2^n} \prod_{j=1}^n \sum_{x_j=0}^1 (-1)^{(s_j \oplus m_j)x_j} \quad (5.7)$$

$$= \frac{1}{2^n} \prod_{j=1}^n (1 + (-1)^{s_j \oplus m_j}) \quad (5.8)$$

Note that the probabilities of measuring any output states $m \neq s$ vanish due to (5.8). Consequently, the amplitudes interfere constructively towards a probability of 1 if every bit of the output state m is equal to s .

5.3 Simon's Algorithm

A well known problem in classical complexity theory is the retrieval of a promised hidden-shift in the pre-image of a boolean function. In 1997, Dan Simon [Sim97] proposed a quantum period finding algorithm with surprising consequences for cryptography, an algorithm in which the secret shift is determined using only linear amounts of queries to the oracle. Recently, Kaplan et al. [KLLP16] showed how to break symmetric-key encryption schemes based on *block-ciphers* using *Simon's algorithm*. In the following, we state the original problem and discuss the relevant algorithm:

Simon's Problem:

Given a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and the promise that there exists $s \in \{0, 1\}^n$ such that, for any pair $(x, y) \in \{0, 1\}^n$, it holds that $f(x) = f(y)$ if and only if $x = y$ or $x \oplus y = s$, the goal is to find s .

The best known classical probabilistic algorithm for collision finding is known to have a complexity of $\Omega(2^{n/2})$ in the number of its queries. Simon's algorithm, the quantum algorithm for solving the above problem, features a query complexity of $O(n)$ [Sim97].

Algorithm 3 Simon's Algorithm

Input: A quantum black box oracle \mathcal{O}_f for a boolean function f with the promise of a hidden-shift $s \in \{0, 1\}^n$. The task is to determine s .

Output: The hidden shift $s \in \{0, 1\}^n$ after $O(n)$ many queries.

Procedure:

1. Initialize $2n$ -qubits and apply the Hadamard gate $H^{\otimes n}$ onto the first n registers:

$$\longrightarrow \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle |0\rangle$$

2. Query the quantum oracle:

$$\longrightarrow \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle |f(x)\rangle$$

3. Perform a measurement onto the second half of n registers, resulting in an outcome $f(z)$ and causing the remaining registers to collapse to the state:

$$\longrightarrow \frac{1}{\sqrt{2^n}} (|z\rangle + |z \oplus s\rangle)$$

4. Apply another Hadamard gate $H^{\otimes n}$ onto the first register:

$$\longrightarrow \frac{1}{\sqrt{2}} \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} (-1)^{y \cdot z} (1 + (-1)^{y \cdot s}) |y\rangle$$

5. Measure the entire output state. Note that, since vectors with the property $s \cdot y = 1$ have an amplitude of 0, a measurement in the computational basis results in a random y that is orthogonal with respect to s . Repeat the previous steps for enough samples of y in order to find s by Gaussian elimination.

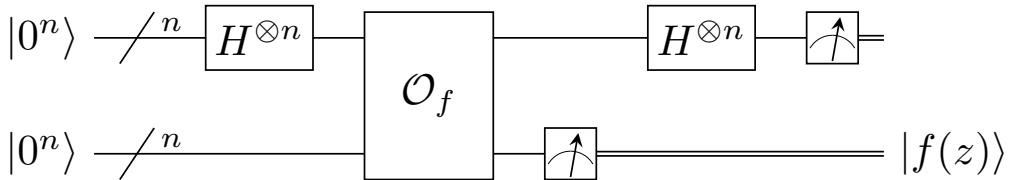


Figure 11: A quantum circuit for Simon's problem. The secret string is determined after $O(n)$ many repetitions given enough linearly independent measurement outcomes.

Typically, after repeating the procedure in the order of only $O(n)$ queries, the full set of basis vectors can be found with high probability. In the case that there are multiple collisions due to more than one secret shift, that is, for any x other than $f(x) = f(x \oplus s)$, Kaplan et al. [KLLP16] provide theorems on the trade-off between the number of repetitions and the success probability of the algorithm. In Chapter 10, we discuss the performance of recent implementations of Simon’s algorithm on two different quantum computing architectures.

6 The Quantum Fourier Transform

The quantum Fourier transform (QFT) is arguably the most important and most widely used tool in the design of quantum algorithms. Many of the known algorithms, such as the Deutsch-Josza algorithm, Shor’s factoring algorithm, quantum phase estimation, quantum order finding, Simon’s algorithm or the Bernstein-Vazirani algorithm rely substantially on its use. In the previous section, we encountered the Hadamard transform as a single qubit gate H that performs the following operation:

$$H|x\rangle = \frac{1}{\sqrt{2}} \sum_{y=0}^1 (-1)^{x \cdot y} |y\rangle. \quad (6.1)$$

Implicitly, we have already encountered the single-qubit quantum Fourier transform of order $n = 1$. In fact, the Hadamard transform can be thought of as a QFT over the group $\mathbb{Z}/2\mathbb{Z}$ that takes single qubits in the computational basis and maps them to the Hadamard basis. The underlying principle of the Fourier transform already starts to show, namely, that the QFT acts as a change of basis in which the amplitudes of individual states of the computational basis are *related* to the amplitudes of the entire computational space. The amplitudes of the transformed states are the so-called *characters* of the Fourier transform. In this chapter, our goal is to introduce a general *qudit* extension of the Hadamard transform, the QFT on the cyclic group $\mathbb{Z}/q\mathbb{Z}$, where q is any integer. By extension, we also obtain the QFT over any finite Abelian group. More importantly, we discuss how the QFT can be efficiently implemented on a quantum computer. Historically, the earliest non-trivial variant known to be efficiently computable on a quantum computer is the Fourier transform over the group $\mathbb{Z}/2^n\mathbb{Z}$, due to Deutsch [Deu85]. The general variant of the quantum Fourier transform where q is an arbitrary integer is less common, but also mentioned as a side note in [NC10]. Kitaev [Kit95] was the first to generalize this variant of the Fourier transform using quantum phase estimation. Evidently, this breakthrough immediately led to the generalization over finite Abelian groups, as we will discuss in detail in the next section. In both cases, we review efficient quantum circuit implementations and thus provide the basis for addressing the QFT in the extended Bernstein-Vazirani problem.

6.1 The Quantum Fourier Transform over Finite Abelian Groups

The Fourier transform can be defined on arbitrary groups and we can extend the same principle of basis change into the language of groups and group algebras. For further reading on the Fourier transform on groups, we refer to the survey [CD10] or the supplementary chapters in [NC10]. In this thesis, we concern ourselves with the case of finite Abelian groups.

Consider a finite Abelian group $(G, *)$ of order $|G| = N$ and let $\mathbb{C}G = \text{span}\{|g\rangle : g \in G\}$ be the associated group algebra of G over \mathbb{C} . Each element $|x\rangle \in \mathbb{C}G$ can be uniquely expressed as a linear combination of basis vectors and complex coefficients:

$$|x\rangle = \sum_{g \in G} \alpha_g |g\rangle, \quad \text{where } \alpha_g \in \mathbb{C}. \quad (6.2)$$

Since G is finite Abelian, there exists a basis $\hat{G} \subseteq \mathbb{C}G$ of dimension $|\hat{G}| = N$ that consists of N distinct irreducible characters $\chi \in \hat{G}$, where $\chi : G \rightarrow \mathbb{C}$ and $\chi(a * b) = \chi(a) \cdot \chi(b)$ ([CD10], Appendix B). Moreover, if $\chi, \chi' \in \hat{G}$ are two irreducible characters, then:

$$\sum_{x \in G} \chi(x) * \overline{\chi'(x)} = N \delta_{\chi, \chi'}. \quad (6.3)$$

Definition 6.1. Let $(G, *)$ be a finite Abelian group of order $|G| = N$ and let \hat{G} be the basis containing the set of N distinct characters $\chi : G \rightarrow \mathbb{C}$.

The quantum Fourier transform \mathcal{F}_G on the group G is defined as the operation:

$$|x\rangle \rightarrow \frac{1}{\sqrt{N}} \sum_{\chi \in \hat{G}} \chi(x) |\chi\rangle. \quad (6.4)$$

Consider now the case where G is given by the group $G = (\mathbb{Z}/N\mathbb{Z}, +)$ and N is any integer. In this case, the irreducible characters $\chi : \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$ are given precisely by the primitive N^{th} roots of unity $\omega_N = e^{\frac{2\pi i}{N}}$. Thus, for every $y \in \mathbb{Z}/N\mathbb{Z}$, the character $\chi_y(x) = \omega_N^{xy}$ is uniquely determined. By choosing an orthonormal basis $|0\rangle, |1\rangle, \dots, |N\rangle$ of \hat{G} in Fourier space, we can identify (6.4) as:

$$|x\rangle \rightarrow \frac{1}{\sqrt{N}} \sum_{y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{xy} |y\rangle. \quad (6.5)$$

Furthermore, we can associate the Fourier transform \mathcal{F}_G with the operator:

$$\mathcal{F}_G = \frac{1}{\sqrt{N}} \sum_{x, y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{xy} |y\rangle \langle x|. \quad (6.6)$$

Let us briefly state the following fact, analogous to as in Eq.(6.3), regarding orthogonality of the roots of unity:

Proposition 6.2.

$$\sum_{y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{xy} \omega_N^{-x'y} = N \delta_{x, x'}. \quad (6.7)$$

Proof. Consider first the case when $x = x'$. Then, for all $y \in \mathbb{Z}/N\mathbb{Z}$, we find $\omega_N^{(x-x')y} = 1$ and the above sum clearly adds up to N . If $x \neq x'$, we can apply the partial sum formula of the geometric series and compute:

$$1 + \omega_N^{x-x'} + (\omega_N^{x-x'})^2 + \dots + (\omega_N^{x-x'})^{N-1} = \frac{1 - (\omega_N^{x-x'})^N}{1 - \omega_N^{x-x'}} = 0. \quad (6.8)$$

□

We can apply Proposition 6.2 in order to show that \mathcal{F}_G is indeed a well defined unitary operation:

$$\mathcal{F}_G \mathcal{F}_G^\dagger = \frac{1}{N} \sum_{x,y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{x \cdot y} |x\rangle \langle y| \sum_{x',y' \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{-y' \cdot x'} |y'\rangle \langle x'| \quad (6.9)$$

$$= \frac{1}{N} \sum_{x,x',y,y' \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{x \cdot y - x' \cdot y'} \delta_{y,y'} |x\rangle \langle x'| \quad (6.10)$$

$$= \frac{1}{N} \sum_{x,x',y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{(x-x') \cdot y} |x\rangle \langle x'| = \sum_{x,x' \in \mathbb{Z}/N\mathbb{Z}} \delta_{x,x'} |x\rangle \langle x'| = \mathbb{1}. \quad (6.11)$$

According to the *fundamental classification of finite abelian groups* [CD10], any finite Abelian group G is structurally equivalent, i.e. isomorphic, to a direct product of cyclic factors whose orders are prime powers. Let $|G| = N$ and let $N = p_1^{r_1} \dots p_k^{r_k}$ be the unique prime factorization of N , then:

$$G \cong \mathbb{Z}/p_1^{r_1}\mathbb{Z} \times \dots \times \mathbb{Z}/p_k^{r_k}\mathbb{Z}. \quad (6.12)$$

Moreover, the basis \hat{G} of irreducible characters is given by products of irreducible characters of the respective factors in (6.12). Consequently, following [CD10], the quantum Fourier transform on finite Abelian groups G is given by:

$$\mathcal{F}_G = \mathcal{F}_{\mathbb{Z}/p_1^{r_1}\mathbb{Z}} \otimes \dots \otimes \mathcal{F}_{\mathbb{Z}/p_k^{r_k}\mathbb{Z}}. \quad (6.13)$$

For example, let $G = ((\mathbb{Z}/N\mathbb{Z})^n, +)$. Then, for any $y \in (\mathbb{Z}/N\mathbb{Z})^n$, we can associate a unique irreducible character $\chi_y : \mathbb{Z}/N\mathbb{Z} \longrightarrow \mathbb{C}$ such that for all $x \in (\mathbb{Z}/N\mathbb{Z})^n$:

$$\chi_y(x) = \chi_y(x_1) \dots \chi_y(x_n) = \omega_N^{x_1 \cdot y_1 + \dots + x_n \cdot y_n}. \quad (6.14)$$

Hence, the quantum Fourier transform \mathcal{F}_G over the group $G = (\mathbb{Z}/N\mathbb{Z})^n$ is given by:

$$|x_1\rangle |x_2\rangle \dots |x_n\rangle \longrightarrow \frac{1}{\sqrt{N^n}} \sum_{y \in (\mathbb{Z}/N\mathbb{Z})^n} \omega_N^{x_1 \cdot y_1 + \dots + x_n \cdot y_n} |y_1\rangle |y_2\rangle \dots |y_n\rangle, \quad (6.15)$$

where, from now on, $\langle x, y \rangle = x_1 \cdot y_1 + \dots + x_n \cdot y_n$. In the following sections, we consider different variants of groups G and derive a quantum circuit implementation that realizes the corresponding Fourier transformations. Before we give efficient circuit implementations, we require additional remarks.

Let G be the group $G = (\mathbb{Z}/N\mathbb{Z}, +)$ and consider the *shift operator* $U(1)$ that performs the following operation:

$$U(1) : |x\rangle \longrightarrow |x+1\rangle, \quad (6.16)$$

where $x+1$ is the cyclic addition *mod* N . We can verify that $U(1) = \sum_{x \in \mathbb{Z}/N\mathbb{Z}} |x+1\rangle \langle x|$ is unitary, since:

$$U(1)U(1)^\dagger = \sum_{x \in \mathbb{Z}/N\mathbb{Z}} |x+1\rangle \langle x| \sum_{x' \in \mathbb{Z}/N\mathbb{Z}} |x'\rangle \langle x'+1| \quad (6.17)$$

$$= \sum_{x, x' \in \mathbb{Z}/N\mathbb{Z}} |x+1\rangle \langle x| |x'\rangle \langle x'+1| = \sum_{x \in \mathbb{Z}/N\mathbb{Z}} |x+1\rangle \langle x+1| = \mathbb{1}. \quad (6.18)$$

Therefore, we can prove the following statement that connects our previous discussion on the quantum Fourier transform with the shift operator, as appeared in the work of Kitaev [Kit95]:

Proposition 6.3. *The shift operator $U(1)$ is diagonal in the Fourier basis:*

$$\mathcal{F}_G U(1) \mathcal{F}_G^\dagger = \sum_{y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^y |y\rangle \langle y|. \quad (6.19)$$

Proof. Let the operators be represented by:

$$\begin{aligned} \mathcal{F}_G &= \frac{1}{\sqrt{N}} \sum_{x, y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{x \cdot y} |y\rangle \langle x|, \quad U(1) = \sum_{z \in \mathbb{Z}/N\mathbb{Z}} |z+1\rangle \langle z| \quad \text{and} \\ \mathcal{F}_G^\dagger &= \frac{1}{\sqrt{N}} \sum_{x', y' \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{-y' \cdot x'} |y'\rangle \langle x'|. \end{aligned}$$

Using Proposition 6.2, we compute:

$$\mathcal{F}_G U(1) \mathcal{F}_G^\dagger = \frac{1}{N} \sum_{x, y, x', y' \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{x \cdot y} \omega_N^{-y' \cdot x'} |y\rangle \langle x| |y'+1\rangle \langle x'| \quad (6.20)$$

$$= \frac{1}{N} \sum_{y, x', y' \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{(y'+1) \cdot y} \omega_N^{-y' \cdot x'} |y\rangle \langle x'| \quad (6.21)$$

$$= \sum_{y, x' \in \mathbb{Z}/N\mathbb{Z}} \omega_N^y \left[\frac{1}{N} \sum_{y' \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{y' \cdot y} \omega_N^{-y' \cdot x'} \right] |y\rangle \langle x'| \quad (6.22)$$

$$= \sum_{y, x' \in \mathbb{Z}/N\mathbb{Z}} \omega_N^y \delta_{y, x'} |y\rangle \langle x'| \quad (6.23)$$

$$= \sum_{y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^y |y\rangle \langle y|. \quad (6.24)$$

□

6.2 Efficient Circuit Implementations

One of the earliest efficient circuit implementations was found for the QFT of order $N = 2^n$, as provided by the following theorem:

Theorem 6.4 ([BV93]). *For any integer n and order $N = 2^n$, there exists an efficient quantum circuit that uses $O(n^2)$ elementary gates and performs the quantum Fourier transform on any of the orthonormal basis states orthonormal basis $|0\rangle, |1\rangle, \dots, |N-1\rangle$:*

$$|x\rangle \longrightarrow \frac{1}{\sqrt{2^n}} \sum_{y \in \mathbb{Z}/2^n\mathbb{Z}} \omega_N^{x \cdot y} |y\rangle. \quad (6.25)$$

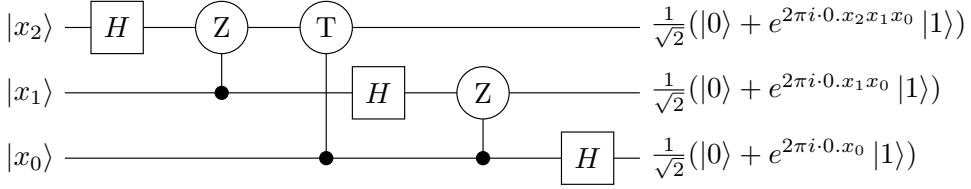


Figure 12: A quantum circuit that performs the QFT for three qubits using the following elementary gates: Hadamard (H), Controlled-Z (Z) and the Controlled- $\pi/4$ Phase-shift (T).

In particular, in the case of $N = 2^n$, we can use a binary representation $x = \sum_{j=0}^{n-1} 2^j x_j$ so that $x = x_1 x_2 \dots x_n$. Moreover, it is also helpful to introduce the binary fraction notation $[0.x_1 \dots x_m] = \sum_{i=1}^m 2^i x_i$. This allows us to write the QFT in terms of a separable product of n -qubits [NC10]:

$$|x\rangle \rightarrow \frac{1}{\sqrt{2^n}} \sum_{y=0}^{2^n-1} \omega_N^{x \cdot y} |y\rangle \quad (6.26)$$

$$= \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} \omega_N^{x \sum_{j=0}^{n-1} 2^j y_j} |y_1\rangle \dots |y_n\rangle \quad (6.27)$$

$$= \frac{1}{\sqrt{2^n}} \bigotimes_{j=0}^{n-1} \sum_{y_j \in \{0,1\}} e^{2\pi i x_j y_j / 2^{n-j}} |y_j\rangle \quad (6.28)$$

$$= \bigotimes_{j=0}^{n-1} \frac{|0\rangle + e^{2\pi i \sum_{k=0}^{n-1} 2^{j+k-n} x_k} |1\rangle}{\sqrt{2}} \quad (6.29)$$

$$(6.30)$$

Next, we would like to extend the QFT onto an arbitrary cyclic group $\mathbb{Z}/N\mathbb{Z}$, by using a technique due to Kitaev [Kit95]. Following [CD10], we can derive this transformation using *quantum phase estimation*, an efficient quantum procedure for the estimation of eigenvalues for a given unitary operator [NC10]. The goal is to perform the QFT over $\mathcal{F}_{\mathbb{Z}/N\mathbb{Z}}$ and map a state $|x\rangle$ into the Fourier basis $|\hat{x}\rangle$, as follows:

$$|x\rangle \longrightarrow |\hat{x}\rangle = \frac{1}{\sqrt{N}} \sum_{y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{x \cdot y} |y\rangle. \quad (6.31)$$

Let us first note that by additionally attaching the input state in Eq.(6.31), it is straightforward to realize the above operation as a two-qubit operation using elementary gates. This can be verified as follows: First prepare the state $|x\rangle|0\rangle$ and create a uniform superposition in the second register:

$$|x\rangle|0\rangle \longrightarrow \frac{1}{\sqrt{N}} \sum_{y \in \mathbb{Z}/N\mathbb{Z}} |x\rangle|y\rangle. \quad (6.32)$$

Consider now applying a controlled phase-shift gate $|x\rangle|y\rangle \longrightarrow \omega_N^{x,y}|x\rangle|y\rangle$. As a result, the output is thus transformed into:

$$\frac{1}{\sqrt{N}} \sum_{y \in \mathbb{Z}/N\mathbb{Z}} |x\rangle|y\rangle \longrightarrow \frac{1}{\sqrt{N}} \sum_{y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{x,y}|x\rangle|y\rangle. \quad (6.33)$$

However, due to entanglement of the registers, straightforward erasure of the first register is not possible. At this point, however, we can make use of the quantum phase estimation procedure that allows us to efficiently approximate the eigenvalues of a given unitary operator with $n = O(\log N)$ bits of precision [NC10]. According to Lemma 6.3, the eigenvalues of the shift operator $U(1)$ are precisely given by the roots of unity ω_N . Thus, we can approximately perform the following unitary operation \mathcal{P} :

$$|\hat{x}\rangle|0\rangle \longrightarrow |\hat{x}\rangle|x\rangle \quad (6.34)$$

By reversing the above operation and applying \mathcal{P}^\dagger , we arrive at the desired outcome:

$$|\hat{x}\rangle|x\rangle \longrightarrow |\hat{x}\rangle|0\rangle. \quad (6.35)$$

Finally, we refer to the following highly efficient realization of the QFT due to Hales and Hallgren:

Theorem 6.5 ([HH00]). *For arbitrary integers N , where $n = \log(N)$, and any $\epsilon > 0$, there exists an efficient quantum circuit that uses $O(n \log \frac{n}{\epsilon} + \log^2 \frac{1}{\epsilon})$ many gates and approximately performs the quantum Fourier transform on orthonormal basis states $|0\rangle, |1\rangle, \dots, |N-1\rangle$ up to a fidelity of ϵ :*

$$|x\rangle \longrightarrow \frac{1}{\sqrt{N}} \sum_{y \in \mathbb{Z}/N\mathbb{Z}} \omega_N^{x,y}|y\rangle. \quad (6.36)$$

7 Quantum Learning Algorithms

Quantum computers can indeed solve certain problems faster than classical computers, as demonstrated in Section 5. The tasks we considered so far all concerned static learning tasks with well-defined entities free of noise and error. As decoherence still poses a major threat to current quantum computing architectures, the promise of successfully running quantum algorithms is still largely dependent on the extent to which fault-tolerant computing is currently realized. The theory of quantum error correction has been crucial in establishing the prospect of fault-tolerant quantum computing in the near future. Currently, noise in quantum computing architectures is regarded as fatal and believed to substantially slow down most quantum improvements over classical algorithms. In this chapter, we review recent work by Cross, Smith and Smolin [CSS14], showing that quantum algorithms can indeed solve certain tasks in the presence of certain classes of noise, much to the contrary of their classical counterparts. We introduce tasks from computational learning theory, such as the *Learning Parity with Noise* (LPN) problem which concerns the decoding of random linear binary codes. The LPN problem is conjectured to be classically intractable, as the best known algorithms require sub-exponential numbers of resources [BKW94]. However, learning in the quantum setting remains easy despite the presence of noise. Furthermore, we discuss its consequences in a related setting in which we consider the LWE problem with quantum samples. To this end, we propose a new generalization of the Bernstein-Vazirani algorithm of Section 5 and present recent results by Grilo and Kerenidis [GK17], demonstrating a successful amplification of the success probability in the presence of noise.

7.1 Computational Learning Theory

We begin by introducing a few basic notions from computational learning theory, following a recent survey on quantum learning theory by Arunachalam and de Wolf [AdW17].

Let us start with a few relevant definitions regarding the objectives of learning.

A *concept class* $\mathfrak{C} = \bigcup_{n \geq 1} C_n$ is a collection of *concepts* (typically Boolean functions) in which each set C_n is to contain all *concepts* $f : \{0, 1\}^n \rightarrow \{0, 1\}$. We consider *learning problems* as a setting in which a *learner* \mathcal{A} , i.e. an algorithm, is given access to either a membership or example oracle for a *target concept* $f \in C_n$ and the task is to then find a *hypothesis* $h \in C_n$ that agrees with the concept f upon some measure of accuracy. In other words, having access to a black box oracle, the goal of the learner is to correctly identify the oracle that corresponds to the target concept. Note that all definitions in this section also translate naturally in the context of a computational space $\mathbb{Z}/q\mathbb{Z}$, where $q \geq 2$ is any integer. Next, we specify variants of learning models, both in the classical, as well as in the quantum setting.

7.1.1 Exact Learning

Definition 7.1 (Classical Exact Learning).

In a classical exact learning model, a learner \mathcal{A} for a concept class C_n is given access to a membership oracle \mathcal{O}_f for a target concept $f \in C_n$ and the task is to find a hypothesis $h \in C_n$ that agrees with the target concept f on all the inputs in $\{0,1\}^n$. Upon input $x \in \{0,1\}^n$, the membership oracle \mathcal{O}_f outputs a label $f(x)$.

We say an efficient algorithm \mathcal{A} is an exact learner for C_n if, for every $f \in C_n$, there exists $\delta > 0$ such that, with probability $1 - \delta$, \mathcal{A} outputs a hypothesis h where for all $x \in \{0,1\}^n : h(x) = f(x)$.

We refer to the query complexity of \mathcal{A} as the maximum number of requests to the membership oracle, over all $f \in C_n$, as well as over the internal randomness needed to achieve the desired success probability of $1 - \delta$.

Definition 7.2 (Quantum Exact Learning).

In a quantum exact learning model, a learner \mathcal{A} for a concept class C_n is given access to a quantum membership oracle \mathcal{O}_f for a target concept $f \in C_n$ and the task is to find a hypothesis $h \in C_n$ that agrees with the target concept f on all the inputs in $\{0,1\}^n$. Upon input $x \in \{0,1\}^n$ and $y \in \{0,1\}$, the membership oracle performs the operation:

$$\mathcal{O}_f : |x\rangle |y\rangle \longrightarrow |x\rangle |y \oplus f(x)\rangle .$$

We say an efficient quantum algorithm \mathcal{A} is an exact learner for C_n if, for every $f \in C_n$, there exists $\delta > 0$ such that, with probability $1 - \delta$, \mathcal{A} outputs a hypothesis h where for all $x \in \{0,1\}^n : h(x) = f(x)$.

Similarly, we now refer to the quantum query complexity of \mathcal{A} as the maximum number of quantum queries to the membership oracle, over all $f \in C_n$, as well as over the internal randomness needed to achieve the desired success probability of $1 - \delta$.

7.1.2 PAC Learning

In this section, we introduce a variant called *probably approximately correct* (PAC) learning, a model in which we consider uniform example oracles contrary to membership oracles. We begin by specifying the learning model in the classical, as well as quantum setting.

Definition 7.3 (Classical PAC Learning).

In a PAC learning model, a learner \mathcal{A} for a concept class C_n is given access to a uniform example oracle $\mathcal{O}_{EX}(f)$ for a target concept $f \in C_n$ and the task is to find a hypothesis $h \in C_n$ that agrees with the target concept f on at least a $1 - \epsilon$ fraction of the inputs in $\{0,1\}^n$.

Upon each query, the example oracle $\mathcal{O}_{EX}(f)$ samples a label $f(x)$ uniformly at random.

We say an algorithm \mathcal{A} is a PAC learner for C_n if, for every $f \in C_n$, there exists an $\epsilon > 0$ and $\delta > 0$ such that, with probability $1 - \delta$, \mathcal{A} outputs a hypothesis h , where:

1. $\Pr_{x \in \{0,1\}^n} [h(x) = f(x)] \geq 1 - \epsilon$.
2. \mathcal{A} runs in time and uses a number of queries that is $\text{poly}(n, 1/\epsilon, 1/\delta)$.

We refer to the query complexity of \mathcal{A} as the maximum number of requests to the example oracle, over all $f \in C_n$, as well as over the internal randomness needed to achieve the desired success probability of $1 - \delta$. The (ϵ, δ) -PAC sample complexity of a concept class C is given by the minimum sample complexity over all (ϵ, δ) -PAC learners for C_n .

Definition 7.4 (Quantum PAC Learning).

In a quantum PAC learning model, a learner \mathcal{A} for a concept class C_n is given access to a quantum example oracle $\mathcal{O}_{EX}(f)$ for a target concept $f \in C_n$ and the task is to find a hypothesis $h \in C_n$ that agrees with the target concept f on at least a $1 - \epsilon$ fraction of the inputs in $\{0, 1\}^n$.

When queried, the example oracle $\mathcal{O}_{EX}(f)$ responds with a quantum state:

$$\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x_1\rangle \dots |x_n\rangle |f(x)\rangle.$$

We say a quantum algorithm \mathcal{A} is a quantum PAC learner for C_n if, for every $f \in C_n$, there exists an $\epsilon > 0$ and $\delta > 0$ such that, with probability $1 - \delta$, \mathcal{A} outputs a hypothesis h , where:

1. $\Pr_{x \in \{0,1\}^n} [h(x) = f(x)] \geq 1 - \epsilon$.
2. \mathcal{A} runs in time and uses a number of queries that is $\text{poly}(n, 1/\epsilon, 1/\delta)$.

We refer to the query complexity of \mathcal{A} as the maximum number of requests (at unit cost) to the example oracle, over all $f \in C_n$, as well as over the internal randomness needed to achieve the desired success probability of $1 - \delta$. The (ϵ, δ) -PAC sample complexity of a concept class C is given by the minimum sample complexity over all (ϵ, δ) -PAC learners for C_n .

Definition 7.5 (PAC Learnable Classes).

We say a concept class $\mathfrak{C} = \bigcup_{n \geq 1} C_n$ is classically (or quantumly) PAC learnable if, given an example oracle for any target concept $f \in \mathfrak{C}$, there exists a PAC algorithm such that, for any $\epsilon, \delta \in (0, 1/2)$, the algorithm

1. outputs an ϵ -approximation h of f with probability $1 - \delta$.
2. runs in time and uses a number of queries that is $\text{poly}(n, 1/\epsilon, 1/\delta)$.

In the next section, we apply these definitions to the *learning party with noise* problem and consider classical, as well as quantum algorithms.

7.2 Learning Parity With Noise

Consider the following well known computational problem resembling a noisy variant of the Bernstein-Vazirani problem in Chapter 5:

Learning Parity With Noise Problem:

Recover the secret $s \in \{0, 1\}^n$ from the class of parity functions $f_s : \{0, 1\}^n \rightarrow \{0, 1\}$ by making queries to a uniform example oracle of Bernoulli noise rate $\eta < 1/2$, where

$$f_s(x) = s_1 \cdot x_1 \oplus \dots \oplus s_n \cdot x_n \bmod 2 = \langle s, x \rangle \bmod 2. \quad (7.1)$$

In the noiseless case, this problem amounts to Gaussian elimination given enough linearly independent samples. Following [CSS14], the probability that n queries to the example oracle \mathcal{O}_{f_s} produce a set of linearly independent examples is given by:

$$(1 - 2^{-n}) \cdot (1 - 2^{-n+1}) \cdots \left(1 - \frac{1}{2}\right) = \prod_{j=0}^{n-1} (1 - 2^{j-n}). \quad (7.2)$$

A simple proof by induction shows that this probability is in fact greater than $1/4$ for any integer $n > 1$. In the noiseless case, the class of parity functions is clearly PAC-learnable. In fact, any algorithm that fails with constant probability less than some $p \in (0, 1)$ can be repeated in the order of $O(\log_{1/p} 1/\delta)$ to reduce the probability of failure below $\delta > 0$. In the case of a noise rate $\eta < 1/2$, it is known that the LPN problem is an average-case version of the NP-hard problem of decoding a linear code, hence the LPN problem is thus classically intractable.

In the quantum setting, this problem remains easy even in the presence of noise, as shown in [CSS14]. Our goal is to first define the LPN problem in a quantum oracle model using uniform quantum samples and show that the LPN problem is quantumly PAC-learnable. First note its resemblance to the Bernstein-Vazirani problem based on queries from a uniform example oracle.

Learning Parity With Noise:

Recover the secret $s \in \{0, 1\}^n$ from the class of parity functions $f_s : \{0, 1\}^n \rightarrow \{0, 1\}$, where $f_s(x) = s_1 \cdot x_1 \oplus \dots \oplus s_n \cdot x_n \bmod 2 = \langle s, x \rangle \bmod 2$, by querying a quantum example oracle $\mathcal{O}_{EX}(f_s, \eta)$ of noise rate $\eta < 1/2$. Upon each query, $\mathcal{O}_{EX}(f_s, \eta)$ outputs uniform quantum sample given by:

$$|\Psi_s\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in (\mathbb{Z}/2\mathbb{Z})^n} |x_1\rangle \dots |x_n\rangle |\langle x, s \rangle \oplus e\rangle,$$

where the error follows $e \sim \text{Bernoulli}(\eta)$.

Let us first treat the problem in the *noiseless* case. Consider the following algorithm, as in [CSS14]:

Algorithm 4 *Quantum Parity Learning*

Input: A quantum example oracle $\mathcal{O}_{EX}(f_s)$ acting as a black box that outputs ideal uniform quantum samples. The task is to determine $s \in \{0,1\}^n$.

Output: The secret string $s \in \{0,1\}^n$ with probability 1/2.

Procedure:

1. Query $\mathcal{O}_{EX}(f_s)$ and receive a uniform quantum example state $|\psi_{f_s}\rangle$, where

$$|\psi_{f_s}\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x_1\rangle |x_2\rangle \dots |x_n\rangle |f_s(x)\rangle$$

2. Perform a Hadamard gate onto all $n + 1$ registers:

$$\longrightarrow \frac{1}{\sqrt{2}} (|0^n\rangle |0\rangle + |s\rangle |1\rangle)$$

3. Measure the entire output state. Read out s if the last register is $|1\rangle$, else output \perp .

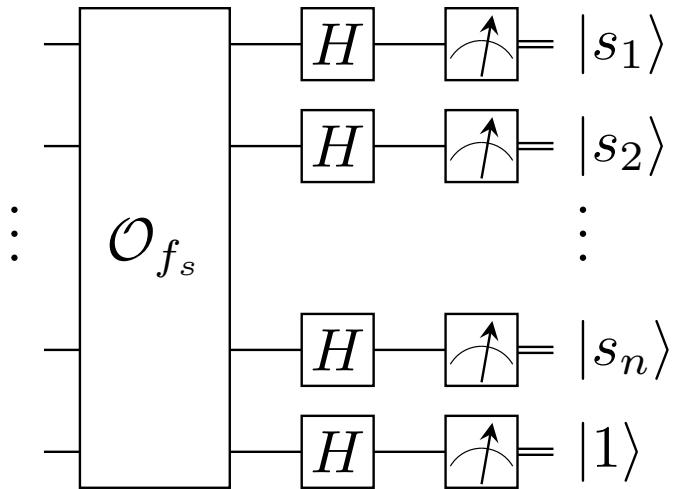


Figure 13: A quantum circuit for the quantum parity learning algorithm. With probability 1/2, the final register is measured to be in the state $|1\rangle$ and the secret string can be read out immediately.

The second step in Algorithm 4 can easily be verified using Proposition 6.2, as follows:

$$\begin{aligned}
H^{\otimes(n+1)} |\psi_{f_s}\rangle &= \frac{1}{\sqrt{2}} \frac{1}{2^n} \sum_{y_{n+1} \in \{0,1\}} \sum_{x,y \in \{0,1\}^n} (-1)^{\langle x,y \rangle} (-1)^{\langle x,s \rangle \cdot y_{n+1}} |y_1\rangle |y_2\rangle \dots |y_n\rangle |y_{n+1}\rangle \\
&= \frac{1}{\sqrt{2}} \left(\frac{1}{2^n} \sum_{x,y \in \{0,1\}^n} (-1)^{\langle x,y \rangle} (-1)^{\langle x,0 \rangle} |y_1\rangle |y_2\rangle \dots |y_n\rangle |0\rangle \right. \\
&\quad \left. + \frac{1}{2^n} \sum_{x,y \in \{0,1\}^n} (-1)^{\langle x,y \rangle} (-1)^{\langle x,s \rangle} |y_1\rangle |y_2\rangle \dots |y_n\rangle |1\rangle \right) \\
&= \frac{1}{\sqrt{2}} \left(\sum_{y \in \{0,1\}^n} \delta_{y,0} |y_1\rangle |y_2\rangle \dots |y_n\rangle |0\rangle + \sum_{y \in \{0,1\}^n} \delta_{y,s} |y_1\rangle |y_2\rangle \dots |y_n\rangle |1\rangle \right) \\
&= \frac{1}{\sqrt{2}} (|0^n\rangle |0\rangle + |s\rangle |1\rangle).
\end{aligned}$$

Let us now consider the LPN problem in its original setting of constant Bernoulli noise rate. Now the learning algorithm is given access to quantum samples that are described as a mixture of both noisy and noiseless samples. Surprisingly, even in this model, the amplitudes interfere constructively as in the previous algorithm after the use of Hadamard gates, as discussed in [CSS14].

Algorithm 5 *Learning Parity With Noise*

Input: A quantum example oracle $\mathcal{O}_{EX}(f_s, \eta)$ acting as a black box that outputs quantum states prone to a parity bit flip error with probability η . The task is to determine s .

Output: The secret string $s \in \{0, 1\}^n$ with probability $1/2$, independent of η .

Procedure:

1. Query $\mathcal{O}_{EX}(f_s, \eta)$ and receive a uniform quantum state $|\psi_{f_s}\rangle$, where $e \sim \text{Bern}(\eta)$:

$$|\psi_{f_s}\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x_1\rangle |x_2\rangle \dots |x_n\rangle |f_s(x) \oplus e\rangle$$

2. Perform a Hadamard gate onto all $n + 1$ registers.

$$\begin{aligned}
&\longrightarrow \frac{1}{\sqrt{2}} (|0^n\rangle |1\rangle + |s\rangle |0\rangle) \quad (\text{with probability } \eta) \\
&\longrightarrow \frac{1}{\sqrt{2}} (|0^n\rangle |0\rangle + |s\rangle |1\rangle) \quad (\text{with probability } 1 - \eta)
\end{aligned}$$

3. Measure the entire output state. Read out any nonzero string, else output \perp .

7.3 Extended Bernstein-Vazirani Algorithm

In this chapter, we introduce a novel *qudit* extension of the well known Bernstein-Vazirani problem in a computational learning setting by considering the problem over the cyclic group $G = (\mathbb{Z}/q\mathbb{Z}, +)$, where q is any integer. In solving this problem, we provide the basis for the quantum LWE problem.

Extended Bernstein-Vazirani Problem:

Recover the secret $s \in (\mathbb{Z}/q\mathbb{Z})^n$ from the class of functions $f_s : (\mathbb{Z}/q\mathbb{Z})^n \rightarrow \mathbb{Z}/q\mathbb{Z}$ by making queries to a uniform example oracle for

$$f_s(x) = s_1 \cdot x_1 + \dots + s_n \cdot x_n \bmod q = \langle s, x \rangle \bmod q, \quad (7.3)$$

where q is any positive integer.

Similar to the noiseless LPN problem, the classical query complexity of the above problem is given by $\Omega(n)$. In the quantum setting, we are given a quantum example oracle \mathcal{O}_{f_s} and the goal is to solve the extended Bernstein-Vazirani problem. In the following, we show that the above problem is exactly learnable by Algorithm 6 and discuss its applications for the LWE problem. Surprisingly, our proposed algorithm achieves a success probability in terms of Euler's totient function $\varphi : \mathbb{N} \rightarrow \mathbb{N}$. Upon input q , the output $\varphi(q)$ is defined as the number of integers k such that $\gcd(k, q) = 1$. Using Euler's product formula, we can also write the probability that Algorithm 6 succeeds as:

$$\frac{\varphi(q)}{q} = \prod_{\text{primes } p|q} \left(1 - \frac{1}{p}\right). \quad (7.4)$$

The curious quotient $\varphi(q)/q$ is of deep importance to number theory and has been studied for many decades. In the 1950ies, Schinzel and Sierpiński proved that $\{\varphi(n)/n : n = 1, 2, \dots\}$ is dense in the interval $(0, 1) \subset \mathbb{R}$, highlighting that the ratio is highly nontrivial. Therefore, it is not possible to find a unique limit as q approaches infinity. Euler's product formula, Eq. (7.4), gives us an intuition on how large $\varphi(q)/q$ is, depending on the prime factorization of q . If q is prime, we observe a simple ratio of $\frac{q-1}{q}$, hence a high probability for Algorithm 6 to succeed. Using a result due to Rosser and Schoenfeld [RS62], we can also bound the success probability of Algorithm 6 in the case where $q > 2$.³

$$\frac{\varphi(q)}{q} > \frac{1}{e^\gamma \log \log(q) + \frac{3}{\log \log(q)}}, \quad (7.5)$$

where $e^\gamma = 1.7810724\dots$ is Euler's constant. For our purposes, this ratio is still constant and the algorithm can be repeated to amplify the success probability as it fails with constant probability less than some $p \in (0, 1)$. Thus, if Algorithm 6 is to succeed after m repetitions with probability $1 - \delta$, we require $O(\log_{1/p} 1/\delta)$ samples and time $\text{poly}(m, n, \log \frac{1}{\delta})$. Surprisingly, the sample complexity is independent of n , whereas the classical query complexity is given by $\Omega(n)$.

³In the case of $q = 2$, the problem is the noiseless variant of the LPN problem from Section 7.2.

Algorithm 6 *Extended Bernstein-Vazirani Algorithm*

Input: A quantum example oracle $\mathcal{O}_{EX}(f_s)$ acting as a black box for the inner product function $f_s(x) = \langle s, x \rangle \bmod q$, where $s \in (\mathbb{Z}/q\mathbb{Z})^n$ is to be determined.

Output: $s \in (\mathbb{Z}/q\mathbb{Z})^n$ with probability $\varphi(q)/q$, where $\varphi(q) = |(\mathbb{Z}/q\mathbb{Z})^\times|$.

Procedure:

1. Query $\mathcal{O}_{EX}(f_s)$ and receive a quantum example:

$$|\Psi_s\rangle = \frac{1}{\sqrt{q^n}} \sum_{x \in (\mathbb{Z}/q\mathbb{Z})^n} |x_1\rangle \dots |x_n\rangle |f_s(x)\rangle.$$

2. Perform the Fourier transform \mathcal{F}_G onto the last register:

$$\frac{1}{\sqrt{q^n}} \frac{1}{\sqrt{q}} \sum_{x \in (\mathbb{Z}/q\mathbb{Z})^n} \sum_{y \in \mathbb{Z}/q\mathbb{Z}} \omega_q^{\langle s, x \rangle \cdot y} |x_1\rangle |x_2\rangle \dots |x_n\rangle |y\rangle$$

3. Measure the last register and obtain a random outcome $k \in \mathbb{Z}/q\mathbb{Z}$:

$$\frac{1}{\sqrt{q^n}} \sum_{x \in (\mathbb{Z}/q\mathbb{Z})^n} \omega_q^{\langle s, x \rangle \cdot k} |x_1\rangle |x_2\rangle \dots |x_n\rangle |k\rangle$$

4. Perform the inverse Fourier transform $\mathcal{F}_G^{\otimes n\dagger}$ onto the first n registers:

$$|ks_1\rangle |ks_2\rangle \dots |ks_n\rangle |k\rangle$$

5. If $\gcd(k, q) = 1$, invert $|ks\rangle$ by multiplying each register by k^{-1} and discard the last register (or else output \perp):

$$|s_1\rangle |s_2\rangle \dots |s_n\rangle$$

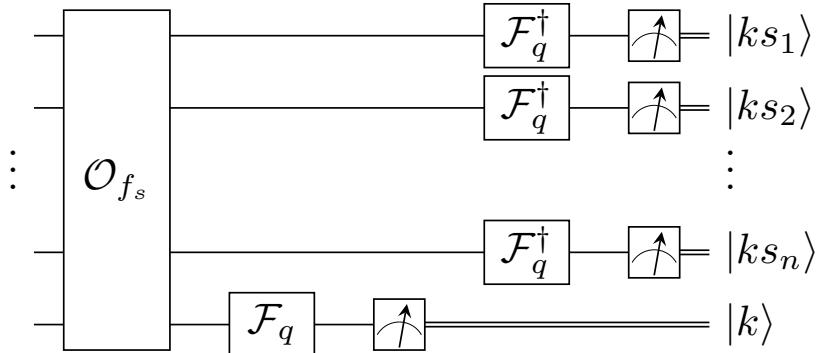


Figure 14: A quantum circuit for the extended Bernstein-Vazirani algorithm.

7.4 Learning With Errors

Finally, we can state the algorithm for the LWE problem with quantum samples. While the Extended-Bernstein-Vazirani algorithm we introduced works for any integer modulus q , Grilo and Kerenidis [GK17] independently proposed a similar algorithm, specifically for the case when q is prime, in order to solve the LWE problem using quantum samples.

Theorem 7.6 ([GK17]). *Let q be a prime in $[2^{n^\gamma}, 2 \cdot 2^{n^\gamma}]$, where $\gamma \in (0, 1)$, and let $\mathcal{O}_{EX}(f_s, \chi)$ be a quantum example oracle for $f_s = \langle s, x \rangle$ that outputs samples*

$$|\Psi_s\rangle = \frac{1}{\sqrt{q^n}} \sum_{x \in (\mathbb{Z}/q\mathbb{Z})^n} |x_1\rangle \dots |x_n\rangle |\langle x, s \rangle + e_x\rangle$$

where the errors e_x are i.i.d. random variables drawn according to $\chi_{\eta, q}$ which is symmetric around 0 and where $\eta = \text{polylog}(q)$.

Let s be the output of the Extended Bernstein-Vazirani algorithm. Then:

1. $\Pr[k = 0 \text{ and output } \perp] = \frac{1}{q}$
2. $\Pr[\text{output} = s \mid \text{output} \neq \perp] \geq \frac{q}{24(q-1)\eta}$
3. $\mathbb{E}[\Pr[\text{output} = s \mid \text{output} \neq \perp]] \leq \frac{1}{q^n}$.

Finally, we also state the result under amplification of the success probability.

Theorem 7.7 ([GK17], Quantum Algorithm for LWE). *For symmetric error distributions $\chi_{\eta, q}$ of noise magnitude $\eta = \text{polylog}(q)$ around 0, the Extended Bernstein-Vazirani algorithm can be amplified to solve LWE _{q, χ} towards a success probability of $1 - \eta$ by requesting $\mathcal{O}(n \log \frac{1}{\eta})$ many samples and running in time $\text{poly}(n, \log \frac{1}{\eta})$.*

8 Blinding of Quantum Algorithms

In the previous chapters, we focused on the exploitation of quantum superposition queries and showed how to achieve speed-ups over classical algorithms. Let us now take a turn towards investigating the limitations of quantum algorithms, in particular in the context of post-quantum cryptography. Our goal is to provide secure cryptographic schemes, even in a setting in which the adversary has quantum access to the encryption or decryption procedure.

Naturally, we investigate the limitations of quantum algorithms in an oracle model. In this setting, a quantum algorithm receives a quantum membership oracle for a given function and is allowed to query the function on a superposition of inputs. We denote quantum access to a function f for a given algorithm \mathcal{A} by adopting the notation $\mathcal{A}^{|f\rangle}$.

8.1 Blinding Lemma

Let us begin by making an important observation: The overlap of two identical n -qubit states remains sufficiently close to unity, even if one of the states is modified at a random location. Therefore, it seems, not even a quantum algorithm receiving output states from an oracle succeeds at determining a random modification in an exponentially large domain with high probability. In fact, we can prove this result as a consequence of the following technical lemma:

Lemma 8.1 (Blinding Lemma).

Let \mathcal{A} be an efficient quantum algorithm making at most $Q = \text{poly}(n)$ queries to an oracle \mathcal{O}_f for a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, where n is a security parameter and $m = \text{poly}(n)$. If $x^* \xleftarrow{\$} \{0, 1\}^n$ is a random location in the domain of f and $s \xleftarrow{\$} \{0, 1\}^m$ is a randomly sampled string, then any output states produced by algorithm \mathcal{A} remain negligibly close in trace distance when replacing the value of $f(x^*)$ with the string s :

$$\mathbb{E}_{s, x^*} [\delta(\mathcal{A}^{\mathcal{O}_f}(1^n), \mathcal{A}^{\mathcal{O}_{f_s}}(1^n))] \leq \frac{2Q}{\sqrt{2^n}}, \quad (8.1)$$

where

$$f_s(x) = \begin{cases} s, & \text{for } x = x^* \\ f(x), & \text{for } x \neq x^*. \end{cases} \quad (8.2)$$

Proof. We can write any QPT algorithm \mathcal{A} with access to an oracle \mathcal{O}_f as a sequence of unitary computations U_0, \dots, U_Q and oracle queries, followed by a measurement. Running $\mathcal{A}(1^n)$ upon an initial state $|\phi_0\rangle$ thus produces an output state according to:

$$|\phi^f\rangle = U_Q \mathcal{O}_f U_{Q-1} \mathcal{O}_f \dots U_1 \mathcal{O}_f U_0 |\phi_0\rangle. \quad (8.3)$$

A final measurement of the output state $|\phi^f\rangle$ through a choice of POVM gives rise to a probability distribution over outcomes according to measurement operators $\mathcal{E} = \{E_i\}$. From Lemma A.1 it follows that, if the trace distance between two output states is bounded by ϵ , then the statistical distance between outcome distributions produced by *any* POVM

over these states is no more than ϵ . Let therefore $x^* \xleftarrow{\$} \{0, 1\}^n$, $s \xleftarrow{\$} \{0, 1\}^m$ and consider $P : \{0, 1\}^n \longrightarrow \{0, 1\}^m$, where

$$P(x) = \begin{cases} f(x^*) \oplus s, & \text{for } x = x^* \\ 0, & \text{for } x \neq x^*. \end{cases} \quad (8.4)$$

If we can show that modifying the oracle functionality from f to $f \oplus P$ results in negligibly close output states, then we can easily conclude that the expected trace distance $\mathbb{E} [\delta(|\phi^f\rangle, |\phi^{f \oplus P}\rangle)]$ must also be negligible. In order to prepare for a hybrid approach, we first show that replacing the functionality of a single oracle query results in statistically close output distributions. To this end, we define the k -th hybrid states as:

$$|\phi_k\rangle = U_Q \mathcal{O}_{f \oplus P} U_{Q-1} \dots \mathcal{O}_{f \oplus P} U_k \mathcal{O}_f \dots \mathcal{O}_f U_0 |\phi_0\rangle \quad (8.5)$$

$$|\phi_k^f\rangle = U_k \mathcal{O}_f U_{k-1} \dots \mathcal{O}_f U_0 |\phi_0\rangle. \quad (8.6)$$

This allows us to bound the total expected distance between the output states as follows:

$$\mathbb{E} [\delta(|\phi^f\rangle, |\phi^{f \oplus P}\rangle)] \leq \mathbb{E} \left[\sum_{k=1}^Q \delta(|\phi_k\rangle, |\phi_{k-1}\rangle) \right] = \sum_{k=1}^Q \mathbb{E} [\delta(|\phi_k\rangle, |\phi_{k-1}\rangle)]. \quad (8.7)$$

Next, we bound two successive hybrids by using invariance of the trace distance with respect to simultaneous unitary transformations:

$$\begin{aligned} \delta(|\phi_k\rangle, |\phi_{k-1}\rangle) &= \delta(U_Q \mathcal{O}_{f \oplus P} \dots \mathcal{O}_{f \oplus P} U_k \mathcal{O}_f \dots \mathcal{O}_f U_0 |\phi_0\rangle, U_Q \mathcal{O}_{f \oplus P} \dots \mathcal{O}_{f \oplus P} U_{k-1} \mathcal{O}_f \dots \mathcal{O}_f U_0 |\phi_0\rangle) \\ &= \delta(\mathcal{O}_f U_{k-1} \dots \mathcal{O}_f U_0 |\phi_0\rangle, \mathcal{O}_{f \oplus P} U_{k-1} \mathcal{O}_f \dots \mathcal{O}_f U_0 |\phi_0\rangle) \\ &= \delta(\mathcal{O}_f |\phi_{k-1}^f\rangle, \mathcal{O}_{f \oplus P} |\phi_{k-1}^f\rangle) \\ &= \delta(|\phi_{k-1}^f\rangle, \mathcal{O}_f \mathcal{O}_{f \oplus P} |\phi_{k-1}^f\rangle) \\ &= \delta(|\phi_{k-1}^f\rangle, \mathcal{O}_P |\phi_{k-1}^f\rangle) \end{aligned}$$

Therefore, we can find the following upper bound for the expected trace distance:

$$\mathbb{E} [\delta(|\phi^f\rangle, |\phi^{f \oplus P}\rangle)] \leq Q \max_{|\psi\rangle} \mathbb{E} [\delta(|\psi\rangle, \mathcal{O}_P |\psi\rangle)] \quad (8.8)$$

$$= Q \max_{|\psi\rangle} \mathbb{E} \left[\sqrt{1 - |\langle \psi | \mathcal{O}_P | \psi \rangle|^2} \right] \quad (8.9)$$

$$\leq Q \max_{|\psi\rangle} \sqrt{1 - \mathbb{E} [|\langle \psi | \mathcal{O}_P | \psi \rangle|^2]}. \quad (8.10)$$

Consider now a projector Π_* onto $\mathbf{supp} \mathcal{O}_P = \text{span}\{ |x^*\rangle |y\rangle \mid y \in \{0, 1\}^n\}$, hence \mathcal{O}_P can now be written as the identity operator, except on the range of Π_* . Using the reverse triangle inequality, we find:

$$|\langle \psi | \mathcal{O}_P | \psi \rangle| = |\langle \psi | \mathcal{O}_P \Pi_* | \psi \rangle + \langle \psi | \mathcal{O}_P (\mathbb{1} - \Pi_*) | \psi \rangle| \quad (8.11)$$

$$\geq -|\langle \psi | \mathcal{O}_P \Pi_* | \psi \rangle| + |1 - \langle \psi | \Pi_* | \psi \rangle| \quad (8.12)$$

$$\geq 1 - 2 \langle \psi | \Pi_* | \psi \rangle. \quad (8.13)$$

Consequently, for any output state $|\phi^f\rangle$ produced by query algorithm \mathcal{A} , we can now bound the expected trace distance as follows:

$$\mathbb{E} \left[\delta(|\phi^f\rangle, |\phi^{f_s}\rangle) \right] \leq Q \max_{|\Psi\rangle} \sqrt{1 - \mathbb{E} [\langle \psi | \mathcal{O}_P | \psi \rangle]^2} \quad (8.14)$$

$$\leq Q \max_{|\Psi\rangle} \sqrt{1 - (1 - 2 \mathbb{E} [\langle \psi | \Pi_* | \psi \rangle])^2} \quad (8.15)$$

$$\leq Q \sqrt{1 - \left(1 - \frac{2}{2^n}\right)^2} \leq \frac{2Q}{\sqrt{2^n}}. \quad (8.16)$$

□

8.2 Relabeling Games

Let us now conclude the blinding lemma from the previous section and formalize the notion of *quantum blindness* towards the class of functions that differ at a random location. To this end, we introduce a new indistinguishability game that allows us to prove the security of our proposed constructions under a quantum chosen-ciphertext attack in the subsequent chapters.

Definition 8.2 (RelabelingGame).

Let \mathcal{O}_f be a quantum oracle for a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and consider the experiment **RelabelingGame** with a QPT algorithm \mathcal{D} , defined as follows:

1. (initial phase) a bit $b \xleftarrow{\$} \{0, 1\}$ and strings $x^* \xleftarrow{\$} \{0, 1\}^n$ and $s \xleftarrow{\$} \{0, 1\}^m$ are generated;
2. (query phase) \mathcal{D} receives oracle access to f , as provided by $\mathcal{O}_f : |x\rangle |y\rangle \rightarrow |x\rangle |y \oplus f(x)\rangle$;
3. (challenge phase) \mathcal{D} continues to have access to \mathcal{O}_f if $b = 0$ only; else \mathcal{D} receives an oracle \mathcal{O}_{f_s} , where

$$f_s(x) = \begin{cases} s, & \text{for } x = r^* \\ f(x), & \text{for } x \neq r^*. \end{cases} \quad (8.17)$$

4. (resolution) \mathcal{D} outputs a bit b' and wins the game if $b' = b$.

We say \mathcal{D} is a *distinguisher* for the **RelabelingGame**, if \mathcal{D} wins with high probability, i.e. with nonnegligible probability better than guessing at random.

Proposition 8.3. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a function. Then any QPT algorithm \mathcal{D} making oracle queries to \mathcal{O}_f wins the **RelabelingGame** with at most negligible probability $\frac{1}{2} + \epsilon(n)$.

Proof. Our goal is to use the blinding argument from Lemma 8.1 and show that any measurement of the output states generated by \mathcal{D} results in negligibly close output distributions, irrespective of which of the oracles is being used by the challenger.

The view of \mathcal{D} is the following: Upon an initial state $|\Psi_0\rangle$, $\mathcal{D}(1^n)$ performs a number of Q unitary computations U_1, U_2, \dots, U_Q , as well as alternating queries to \mathcal{O}_f prior to the challenge phase:

$$|\Psi^f\rangle = U_Q \mathcal{O}_f U_{Q-1} \mathcal{O}_f U_{Q-2} \dots U_2 \mathcal{O}_f U_1 \mathcal{O}_f U_0 |\Psi_0\rangle. \quad (8.18)$$

During the challenge phase, \mathcal{D} is given oracle access to \mathcal{O}_φ , where $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and the goal is to determine whether $\varphi = f$ or $\varphi = f_s$, for some $x^* \xleftarrow{\$} \{0, 1\}^n$ and $s \xleftarrow{\$} \{0, 1\}^m$. Thus, \mathcal{D} proceeds in the challenge query phase (denoted by \uparrow) and generates a quantum state after a total of T queries:

$$|\Psi_\varphi^f\rangle = U_T \mathcal{O}_\varphi U_{T-1} \mathcal{O}_\varphi U_{T-2} \dots U_{Q+2} \mathcal{O}_\varphi U_{Q+1} \uparrow U_Q \mathcal{O}_f U_{Q-1} \dots U_2 \mathcal{O}_f U_1 \mathcal{O}_f U_0 |\Psi_0\rangle. \quad (8.19)$$

According to Lemma 8.1, any quantum query algorithm produces negligibly close output states if the underlying functions differ at a single location. Consequently, any output states produced by \mathcal{D} must also lie within negligible trace distance, as we can bound the distance between the output states as follows:

$$\mathbb{E}_{s, x^*} \left[\delta(|\Psi_f^f\rangle, |\Psi_{f_s}^f\rangle) \right] \leq \mathbb{E}_{s, x^*} \left[\delta(|\Psi_f^f\rangle, |\Psi_{f_s}^{f_s}\rangle) \right] \leq \frac{2T}{\sqrt{2^n}}. \quad (8.20)$$

Using Markov's inequality, we conclude that for a negligible distance $2^{-n/4}$:

$$\Pr_{s, r^* \xleftarrow{\$} \{0, 1\}^n} \left[\delta(|\Psi_f^f\rangle, |\Psi_{f_s}^f\rangle) \geq 2^{-n/4} \right] \leq 2^{n/4} \mathbb{E}_{s, r^*} \left[\delta(|\Psi_f^f\rangle, |\Psi_{f_s}^{f_s}\rangle) \right] \leq 2^{-n/4+1} T.$$

Hence, from Lemma A.1, it follows that any POVM measurement of the final output states of \mathcal{D} reveals at most negligibly close outcome distributions. Finally, we have:

$$\left| \Pr_{s, r^* \xleftarrow{\$} \{0, 1\}^n} [\mathcal{D}^{|f_s\rangle}(1^n) = 1] - \Pr [\mathcal{D}^{|f\rangle}(1^n) = 1] \right| \leq 2^{-n/4+1} T = \epsilon(n). \quad (8.21)$$

□

The **RelabelingGame** provides us with an important limitation of all quantum query algorithms, in particular when proving the security of our proposed constructions from the next chapter. As a direct consequence of blinding, we can prove the indistinguishability of several hybrid games in the next chapter on post-quantum cryptography.

9 Post-Quantum Cryptography

Let us now extend the security notions behind chosen-ciphertext attacks from Chapter 3.3 to a world of quantum computers. In particular, we consider adversaries who receive quantum oracle access to both encryption and decryption at various times during the security game. While the case of quantum CCA2-security has already been introduced in [BZ13], we investigate a less powerful model by considering *non-adaptive quantum chosen-ciphertext attacks*. In this security notion of quantum CCA1, an adversary is given quantum superposition access to both encryption and decryption prior to the challenge phase, followed by a final phase of adaptive challenge access to the encryption oracle.

Our goal is to exploit the blindness of quantum query algorithms towards the class of functions that only differ at a single location in order to provide secure constructions under a non-adaptive quantum chosen-ciphertext attack. To this end, we first define both the IND-QCCA1, a close variant of IND-QCCA1 as well as the SEM-QCCA1 security game, and then propose schemes based on quantum-secure pseudorandom functions and permutations that fulfill our definitions.

9.1 Security Under Non-adaptive Quantum Chosen-Ciphertext Attacks

In this section, we extend the definitions from Chapter 3.3 and introduce notions of security in the context of quantum adversaries. In providing a quantum encryption oracle Enc_k , each query is answered by choosing a randomness and encrypting each message in the superposition from the r -family of unitary operations such that:

$$\text{Enc} : \sum_{m,c} \alpha_{m,c} |m\rangle |c\rangle \longrightarrow \sum_{m,c} \alpha_{m,c} |m\rangle |c \oplus \text{Enc}_k(m; r)\rangle \quad (9.1)$$

Typically, we consider the case of sampling a randomness $r \xleftarrow{\$} \{0,1\}^n$ of equal length to a message space, where $m \in \{0,1\}^n$. Moreover, we consider the quantum decryption oracle Dec_k to be deterministic, hence each oracle query is answered upon a superposition of ciphers as follows:

$$\text{Dec} : \sum_{c,s} \beta_{c,s} |c\rangle |s\rangle \longrightarrow \sum_{c,s} \beta_{c,s} |c\rangle |s \oplus \text{Dec}_k(c)\rangle \quad (9.2)$$

Note that, since Enc_k and Dec_k are required to be PPT algorithms provided by the underlying symmetric-key encryption scheme, both (9.1) and (9.2) correspond to efficient and reversible quantum operations.

9.1.1 Indistinguishability

We begin by first introducing a notion of indistinguishability in the context of a quantum chosen-ciphertext attacks.

Definition 9.1 (IND-QCCA1 Security).

Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme and consider the INDGame between a QPT adversary \mathcal{A} and challenger \mathcal{C} , defined as follows:

1. (initial phase) On input 1^n , \mathcal{C} generates a key $k \leftarrow \text{KeyGen}(1^n)$ and a bit $b \leftarrow \{0, 1\}$;
2. (pre-challenge phase) \mathcal{A} receives oracles Enc_k and Dec_k , then sends (m_0, m_1) to \mathcal{C} ;
3. (challenge phase) \mathcal{C} replies with $\text{Enc}_k(m_b)$ and \mathcal{A} receives an oracle for Enc_k only;
4. (resolution phase) \mathcal{A} outputs a bit b' , and wins if $b = b'$.

We say Π has indistinguishable encryptions under non-adaptive quantum chosen-ciphertext attack (or is IND-QCCA1-secure) if, for every QPT \mathcal{A} , there exists a negligible function $\epsilon(n)$ such that: $\Pr[\mathcal{A} \text{ wins INDGame}] \leq 1/2 + \epsilon(n)$.

9.1.2 Decisional Indistinguishability

In a similar manner, we can define a decisional variant of the indistinguishability security game in which the goal of the adversary is to decide whether the challenge corresponds to an encryption of a previously selected message, or an encryption of a random message. By means of an elementary simulation proof, we can easily observe an equivalence in security of the two definitions.

Definition 9.2 (DecIND-QCCA1 Security).

Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme and consider the DecINDGame between a QPT adversary \mathcal{A} and challenger \mathcal{C} , defined as follows:

1. (initial phase) On input 1^n , \mathcal{C} generates a key $k \leftarrow \text{KeyGen}(1^n)$ and a bit $b \leftarrow \{0, 1\}$;
2. (pre-challenge phase) \mathcal{A} receives oracles Enc_k and Dec_k , then sends m to \mathcal{C} ;
3. (challenge phase) \mathcal{C} replies with $\text{Enc}_k(m)$, if $b = 0$, or else with $\text{Enc}_k(u)$ upon a uniformly random message. Then, \mathcal{A} receives an oracle for Enc_k only;
4. (resolution phase) \mathcal{A} outputs a bit b' , and wins if $b = b'$.

We say that Π has decisionally indistinguishable encryptions under non-adaptive quantum chosen-ciphertext attack (or is DecIND-QCCA1-secure) if, for every QPT \mathcal{A} , we have $\Pr[\mathcal{A} \text{ wins DecINDGame}] \leq 1/2 + \epsilon(n)$.

Proposition 9.3. Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme. Then, Π is IND-QCCA1-secure if and only if Π is DecIND-QCCA1-secure.

9.1.3 Semantic Security

In this section, we first introduce semantic security under a QCCA1 learning phase and then prove an important equivalence between our notion of indistinguishability and semantic security.

Definition 9.4 (SEM-QCCA1). Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be an encryption scheme, and consider the experiment **SEMGame** with a QPT \mathcal{A} , defined as follows.

1. (initial phase) A key $k \leftarrow \text{KeyGen}(1^n)$ and bit $b \leftarrow \{0, 1\}$ are generated;
2. (pre-challenge phase) \mathcal{A} receives access to oracles Enc_k and Dec_k , then outputs a classical challenge template consisting of (Samp, h, f) ;
3. (challenge phase) A plaintext $m \leftarrow \text{Samp}$ is generated; \mathcal{A} receives $h(m)$ and an oracle for Enc_k only; if $b = 1$, \mathcal{A} also receives $\text{Enc}_k(m)$.
4. (resolution) \mathcal{A} outputs a string s , and wins if $s = f(m)$.

We say Π is semantically secure under non-adaptive quantum chosen ciphertext attack (or is SEM-QCCA1) if, for every QPT \mathcal{A} , there exists a QPT \mathcal{S} such that the challenge templates output by \mathcal{A} and \mathcal{S} are identically distributed, and there exists a negligible function $\epsilon(n)$ such that:

$$\left| \Pr_{k \leftarrow \mathcal{K}} [\mathcal{A}(1^n, \text{Enc}_k(m), h(m)) = f(m)] - \Pr[\mathcal{S}(1^n, |m|, h(m)) = f(m)] \right| \leq \epsilon(n),$$

where, in both cases, the probability is taken over plaintexts $m \leftarrow \text{Samp}$.

9.1.4 Equivalence of Indistinguishability and Semantic Security

Let us now show the equivalence of the notions we introduced in this chapter.

Theorem 9.5. Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme. Then, Π is IND-QCCA1-secure if and only if Π is SEM-QCCA1-secure.

Proof. Suppose Π is IND-QCCA1-secure, i.e. has indistinguishable encryptions. Let \mathcal{A} be a QPT algorithm against SEM that receives a challenge $\text{Enc}_k(m)$. Define a QPT simulator \mathcal{S} that also challenges SEM but simply runs \mathcal{A} as a subroutine as follows: Instead of receiving $\text{Enc}_k(m)$ during the SEM challenge, \mathcal{S} relies only on the side information $h(m)$, in particular the plaintext length $|m|$, and simulates \mathcal{A} 's encryption and decryption oracles by making use of its own QCCA1 learning phase. At the challenge phase, \mathcal{S} simply encrypts the string $1^{|m|}$ and forwards $\text{Enc}_k(1^{|m|})$ to \mathcal{A} . After another emulated learning phase, \mathcal{S} finally outputs the same target $f(m)$ that \mathcal{A} outputs. Since Π has indistinguishable encryptions by assumption, \mathcal{A} 's success probability must be negligibly close to the original SEM game of \mathcal{S} .

Now, suppose Π is not IND-QCCA1-secure, hence there exists a QPT distinguisher \mathcal{A} against IND-QCCA1-security. This allows us to build a distinguisher \mathcal{D} running \mathcal{A} as a subroutine against the SEM security game as follows: By using its oracles from the QCCA1

learning phases, \mathcal{D} simulates the IND-QCCA1 security game of \mathcal{A} by simply forwarding all queries to its own oracles. At the QIND challenge phase, \mathcal{A} prepares two messages (m_0, m_1) and presents them to \mathcal{D} . Then, \mathcal{D} prepares a SEM challenge template (U, h, f) , where U describes the uniform distribution over plaintexts $\{m_0, m_1\}$, the side information is given by the length of the messages and where the target function $f(m)$ concerns the function that distinguishes between m_0 and m_1 , i.e. $f(m_0) = 0$ and $f(m_1) = 1$. Using this SEM template, \mathcal{D} receives a ciphertext $\text{Enc}_k(m)$ and presents it to \mathcal{A} as an IND-QCCA1 challenge. Finally, \mathcal{D} simply outputs whatever target bit \mathcal{A} outputs. By assumption, \mathcal{A} succeeds with nonnegligible probability and therefore \mathcal{D} breaks the SEM-QCCA1 security game. \square

9.2 Quantum-secure Pseudorandom Functions and Permutations

In order to find constructions for quantum-secure symmetric-key cryptography, we require the use of appropriate building blocks. Let us now extend the concept of secure pseudorandom functions from Chapter 3.4 to a setting in which an adversary in possession of a quantum computer can query the function on a superposition of inputs. Remarkably, one can find quantum-secure constructions for pseudorandom functions that are secure in this model. Moreover, as shown by Zhandry [Zha12], one can construct quantum-secure PRF's either from the assumption that LWE is hard with classical samples or from the existence of quantum one-way functions.

Definition 9.6 (Quantum-secure Pseudorandom Function).

Let $\{f_k\}_{k \in \mathcal{K}}$ be a family of pseudorandom functions on a key-space \mathcal{K} , a domain \mathcal{X} and a range \mathcal{Y} . $\text{QPRF} = \{f_k\}_{k \in \mathcal{K}}$ is a function family of quantum-secure pseudorandom functions if (for every choice of key k) f_k looks indistinguishable from a perfectly random function, hence if, for every QPT distinguisher \mathcal{D} , there exists a negligible function $\epsilon(n)$ such that:

$$\left| \Pr_{k \xleftarrow{\$} \mathcal{K}} [\mathcal{D}^{|f_k\rangle}(1^n) = 1] - \Pr_{f \xleftarrow{\$} \{\mathcal{F}: \mathcal{X} \rightarrow \mathcal{Y}\}} [\mathcal{D}^{|f\rangle}(1^n) = 1] \right| \leq \epsilon(n) \quad (9.3)$$

A special variant of pseudorandom functions is realized in keyed permutations on strings, which can be obtained directly from purely standard assumptions on the existence of one-way functions [Zha16]. We distinguish between quantum-secure and strong quantum-secure pseudorandom permutations.

Definition 9.7 (Quantum-secure Pseudorandom Permutation).

Let $\{\pi_k\}_{k \in \mathcal{K}}$ be a family of keyed permutations operating on bit strings of domain \mathcal{X} upon a key-space \mathcal{K} . $\text{QPRP} = \{\pi_k\}_{k \in \mathcal{K}}$ is a function family of quantum-secure pseudorandom permutations if (for every choice of key k) π_k looks indistinguishable from a random permutation on strings in \mathcal{X} , hence if, for every QPT distinguisher \mathcal{D} , there exists a negligible function $\epsilon(n)$ such that:

$$\left| \Pr_{k \xleftarrow{\$} \mathcal{K}} [\mathcal{D}^{|f_k\rangle}(1^n) = 1] - \Pr_{\pi \xleftarrow{\$} S_{\mathcal{X}}} [\mathcal{D}^{|f\rangle}(1^n) = 1] \right| \leq \epsilon(n), \quad (9.4)$$

where $S_{\mathcal{X}}$ denotes the set of all permutations over strings in \mathcal{X} .

In order to realize security under a quantum chosen-ciphertext attack, we require an even stronger notion of pseudorandomness, in particular in the presence of a decryption oracle. Fortunately, there exist constructions for the following standard, as considered in [Zha16].

Definition 9.8 (Strong Quantum-secure Pseudorandom Permutation).

A family of pseudorandom permutations $\{\pi_k\}_{k \in \mathcal{K}}$ is a function family of strong quantum-secure pseudorandom permutations QPRP if both π_k and its inverse permutation π_k^{-1} look indistinguishable from a random permutation, hence if, for every QPT distinguisher \mathcal{D} , there exists a negligible function $\epsilon(n)$ such that:

$$\left| \Pr_{k \xleftarrow{\$} \mathcal{K}} [\mathcal{D}^{|\pi_k\rangle|\pi_k^{-1}\rangle}(1^n) = 1] - \Pr_{\pi \xleftarrow{\$} S_{\mathcal{X}}} [\mathcal{D}^{|\pi\rangle|\pi^{-1}\rangle}(1^n) = 1] \right| \leq \epsilon(n). \quad (9.5)$$

Finally, we provide schemes based on the above building blocks of pseudorandom functions and permutations that are quantumly secure under a quantum-chosen ciphertext attack we introduced in this thesis.

9.3 Secure Constructions

In this section, we prove the post-quantum security of two symmetric-key encryption schemes. For the first scheme, we revisit the PRF scheme in Construction 3.8, as introduced in Chapter 3.4, and show that it is secure under a non-adaptive quantum chosen-ciphertext attack. The intuition is that, once the pseudorandom function is taken to be quantum-secure, the pre-challenge phase reveals at most a polynomial amount of evaluations of the pseudorandom function, despite the presence of a decryption oracle. Note that, in this scheme, the encryption oracle only reveals a single functional evaluation of the PRF at a random location at a time. The decryption oracle, however, can additionally serve as a membership oracle for the function of the underlying encryption scheme. Thus, the adversary can generate superposition queries to a PRF f_k over the entire input space by simply preparing a uniform superposition over the first register and initializing the second register to the all-zero state:

$$\text{Dec}_k : \sum_{r \in \{0,1\}^n} |r\rangle |0\rangle \longrightarrow \sum_{r \in \{0,1\}^n} |r\rangle |f_k(r)\rangle. \quad (9.6)$$

At first sight, however, it is not clear whether being able to generate superpositions gives the adversary additional power during the challenge phase. Therefore, we have to bound the amount of information that quantum query algorithms can learn in a suitable way. Our goal is to show that this advantage is still negligible and that Construction 3.8 has decisionally indistinguishable encryptions, even in the presence of decryption oracles under a non-adaptive quantum chosen-ciphertext attack. Finally, due to the equivalence results from the previous section, our proposed scheme then also satisfies indistinguishability of encryptions and semantic security.

Theorem 9.9. *Let QPRF be a family of quantum-secure pseudorandom functions. Then, $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ from Construction 3.8 is DeclIND-QCCA1 secure.*

Proof. Our goal is to show that any QPT adversary \mathcal{A} wins the DecIND security game with a QCCA1 learning phase with at most negligible probability. To this end, we introduce a sequence of indistinguishable hybrid games until we arrive at a security game in which the challenge is perfectly hidden and the adversary cannot win.

First, we replace the QPRF from the original security game with a perfectly random function. According to QPRF security, we only negligibly affect the overall success probability. In the setting where the adversary is classical, this hybrid typically completes the proof, as the probability that the challenge randomness is being revealed outside the challenge is negligible. However, this does not translate directly to a QPT adversary performing quantum queries, as discussed at the beginning of the chapter.

We proceed with a final hybrid game in which we use that quantum query algorithms are *blind* towards the class of functions that only differ at a single location, as in Lemma 8.1. To this end, we modify the same random function at a single location during the challenge phase in a way that is indistinguishable to the adversary, thus making it impossible to succeed for the remainder of the game.

With this in mind, we consider the following sequence of hybrid games:

GAME 0: In this hybrid, the adversary is playing the standard DecIND security game for the scheme in Construction 3.8. Prior to the challenge, the adversary chooses a message m and is given quantum superposition access to both, the encryption oracle Enc_k , as well as the decryption oracle Dec_k . Upon receiving a challenge cipher c^* , the adversary may perform additional queries to Enc_k only and then decide whether the cipher corresponds to a genuine encryption $(r^*, f_k(r^*) \oplus m)$ or an encryption of a uniformly random string $(r^*, f_k(r^*) \oplus u)$.

GAME 1: Replace the QPRF f_k from the previous scheme with a perfectly random function f throughout the entire security game. The challenge is now to distinguish the pair $(r^*, f(r^*) \oplus m)$ from an encryption of a uniformly random string $(r^*, f(r^*) \oplus u)$.

GAME 2: In this hybrid, the challenger adopts the same random function but now keeps track of all randomness values previously used for encryption. If a collision occurs and an encryption of the challenge ciphertext is to be answered with a randomness previously seen by the adversary, the challenger aborts the game and the adversary wins.

GAME 3: In the final hybrid, the challenger keeps track of all randomness values as before, as well as adopts the same random function f until the challenge phase of the game. Then, at the start of the phase, the challenger replaces the challenge value $f(r^*)$ with a uniformly random value $s \leftarrow \{0, 1\}^n$ and, for the remainder of the game, continues answering encryption queries with the modified function:

$$f_s(x) = \begin{cases} s, & \text{for } x = r^* \\ f(x), & \text{for } x \neq r^*. \end{cases} \quad (9.7)$$

Suppose there exists a QPT adversary \mathcal{A} that wins GAME 0 with nonnegligible probability, hence there exists some polynomial $p(n)$ such that:

$$\Pr[\mathcal{A} \text{ wins DecINDGame}] \geq 1/2 + 1/p(n). \quad (9.8)$$

Our claim is that the same adversary must (up to at most negligible probability) also succeed at GAME 3 in which the challenge is perfectly hidden by means of a hybrid argument.

GAME 0 vs. GAME 1:

Since we assumed the PRF f_k to be quantum-secure, we can replace it with a perfectly random function and only negligibly affect the success probability in Eq.(9.8). Note that an adversary \mathcal{A} that succeeds at GAME 0 but not at GAME 1 (i.e. \mathcal{A} only wins with at most negligible probability), allows us to build an f_k -distinguisher that violates QPRF security. We can verify this fact, as follows: Let \mathcal{D} be the distinguisher that is given quantum oracle access \mathcal{O}_φ to a function $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ whose goal is to output 0, whenever φ is a perfectly random function and to output 1, if $\varphi = f_k$ for some k . \mathcal{D} now proceeds as follows: It simulates the DecIND-QCCA1 security game of \mathcal{A} by responding to quantum queries using oracles Enc_φ and Dec_φ that can be implemented based on the output of \mathcal{O}_φ . If \mathcal{A} wins the DecIND-QCCA1 game and outputs $b' = b$, then \mathcal{D} outputs 1, and outputs 0 otherwise. Note that, if $\varphi = f_k$, then the output of \mathcal{D} when running \mathcal{A} is identical to that of \mathcal{A} in GAME 0. Similarly, if $\varphi = f$, then the output is identical to that of \mathcal{A} in GAME 1. The adversary \mathcal{A} is efficient, hence \mathcal{D} is also efficient by construction. Moreover, since we assumed that \mathcal{A} succeeds with nonnegligible probability in GAME 0, so does the distinguisher \mathcal{D} , and for infinitely many n :

$$\left| \Pr_{k \xleftarrow{\$} \{0, 1\}^n} [\mathcal{D}^{|f_k\rangle}(1^n) = 1] - \Pr_{f \xleftarrow{\$} \{\mathcal{F} : \{0, 1\}^n \rightarrow \{0, 1\}^n\}} [\mathcal{D}^{|f\rangle}(1^n) = 1] \right| \geq 1/p(n) - \epsilon(n).$$

Therefore, our crucial assumption of QPRF security is broken by this distinguisher and we conclude that \mathcal{A} must also win GAME 1 with nonnegligible probability.

GAME 1 vs. GAME 2:

Since the challenger keeps track of all randomness values \mathcal{R} previously used for encryption, the adversary is guaranteed to receive fresh randomness in $\{0, 1\}^n \setminus \mathcal{R}$, else the game is aborted. Moreover, as the probability of sampling the same randomness twice is negligible, the difference between these two games is only negligible with regard to the overall success probability in the security game.

GAME 2 vs. GAME 3:

At first sight, the challenge from GAME 2 looks uniformly random. However, since the adversary has quantum oracle access to the function f , it is possible to generate superpositions over the entire input space. In principle, these states contain quantum information on the challenge value $f(r^*)$ and our goal is to show that this advantage is still negligible. We make use of a blinding argument and argue that, for any quantum adversary playing

DecIND-QCCA1, GAME 2 and GAME 3 are indistinguishable. Thus, up to negligible probability, the view of the adversary is that of GAME 3 in which the challenge is perfectly hidden and he cannot succeed.

Any quantum adversary playing the **DecIND-QCCA1** game can be separated into a quantum query routine $\mathcal{A}_0(1^n)$ prior to the challenge, as well as final routine $\mathcal{A}_1(1^n)$ at the challenge phase. Consequently, this must also hold for the QPT adversary \mathcal{A} from GAME 2, whose routines we can describe as follows: Upon an initial state $|\Psi_0\rangle$, $\mathcal{A}_0(1^n)$ performs Q unitary computations U_1, U_2, \dots, U_Q and alternating queries to the encryption and decryption oracle and outputs:

$$|\Psi^f\rangle = U_Q \mathcal{O}_{Dec} U_{Q-1} \mathcal{O}_{Enc} U_{Q-2} \dots U_2 \mathcal{O}_{Dec} U_1 \mathcal{O}_{Enc} U_0 |\Psi_0\rangle. \quad (9.9)$$

At the challenge phase, $\mathcal{A}_1(1^n)$ continues to perform encryption oracle queries with respect to f , as well as final unitary computations $U_{Q+1}, U_{Q+2}, \dots, U_T$ upon $|\Psi^f\rangle$ and generates:

$$|\Psi_f^f\rangle = U_T \mathcal{O}_{Enc} U_{T-1} \dots U_{Q+2} \mathcal{O}_{Enc} U_{Q+1} |\Psi^f\rangle, \quad (9.10)$$

where T is the total number of queries and unitary computations. Finally, after completing both routines, \mathcal{A} measures the final output state and outputs b' .

Suppose now that the adversary \mathcal{A} succeeds at GAME 2 but not at GAME 3, i.e. \mathcal{A} only wins with at most negligible probability. We argue that the same adversary must then win the **RelabelingGame**, hence must violate Proposition 8.3.

Let \mathcal{D} be the distinguisher playing the **RelabelingGame** and receiving a quantum oracle \mathcal{O}_φ for a function $\varphi : \{0,1\}^n \rightarrow \{0,1\}^n$ whose goal is to output 0, whenever φ is a perfectly random function and to output 1, if $\varphi = f_s$ for some $x^* \xleftarrow{\$} \{0,1\}^n$ and $s \xleftarrow{\$} \{0,1\}^m$. Let \mathcal{D} now proceed as follows: It simulates the **DecIND-QCCA1** security game of \mathcal{A} by responding to quantum queries using oracles Enc_φ and Dec_φ that can be implemented based on the output of \mathcal{O}_φ . If \mathcal{A} wins the **DecIND** game and outputs $b' = b$, then \mathcal{D} outputs 1, and outputs 0 otherwise. Note that, if $\varphi = f$, then the output of \mathcal{D} when running \mathcal{A} is identical to that of \mathcal{A} in GAME 2. Similarly, if $\varphi = f_s$, then the output is identical to that of \mathcal{A} in GAME 3. The adversary \mathcal{A} is efficient, hence \mathcal{D} is also efficient by construction. Moreover, since we assumed that \mathcal{A} succeeds with nonnegligible probability in GAME 2, so does the distinguisher \mathcal{D} , and for infinitely many n :

$$\left| \Pr_{s, r^* \xleftarrow{\$} \{0,1\}^n} [\mathcal{D}^{|f_s\rangle}(1^n) = 1] - \Pr[\mathcal{D}^{|f\rangle}(1^n) = 1] \right| \geq 1/p(n) - \epsilon(n). \quad (9.11)$$

Therefore, we conclude that \mathcal{A} must also win GAME 3 with nonnegligible probability, else we observe a violation of Proposition 8.3.

Note that, in GAME 3, the **DecIND** challenge phase now amounts to the impossible task of distinguishing between two uniformly random strings. Therefore, since the view of \mathcal{A} is indistinguishable from such a challenge of perfect secrecy, we must now arrive at an overall contradiction and conclude that no such QPT algorithm \mathcal{A} with nonnegligible success probability exists.

In summary, we must therefore conclude that for any quantum algorithm \mathcal{A} playing the DeclIND-QCCA1 game:

$$\Pr[\mathcal{A} \text{ wins DeclINDGame}] = \frac{1}{2} + \epsilon(n). \quad (9.12)$$

□

Finally, as a direct consequence of the equivalence results of Section 9.1.4, let us conclude the previous result with the following additional observation:

Corollary 9.10. *Let QPRF be a function family of quantum-secure pseudorandom functions. Then, $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ from Construction 3.8 is both IND-QCCA1-secure, as well as SEM-QCCA1-secure.*

As pointed out classically in [KL15], the PRF scheme is easily malleable by an adversary with adaptive decryption oracle access in a CCA2 learning phase. Therefore, Construction 3.8 is neither secure under classical nor quantum adaptive chosen-ciphertext attacks.

Finally, we provide another symmetric-key encryption scheme based on pseudorandom permutations that recently appeared in work by Gagliardoni et al. [GHS16]. As for the previous construction, we follow a similar proof and introduce indistinguishable hybrids and a blinding argument.

Construction 9.11. *For a security parameter n , let both $\mu = \text{poly}(n)$ and $\tau = \text{poly}(n)$. Upon a key space $\mathcal{K} = \{0, 1\}^{\mu+\tau}$, consider a function family of keyed permutations $\{\pi_k\}_{k \in \mathcal{K}}$ operating on bit strings of length $\mu + \tau$ and define a symmetric-key encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ as follows:*

1. (key generation) KeyGen : on input 1^n , generate a key $k \xleftarrow{\$} \{0, 1\}^n$;
2. (encryption) Enc_k : on message $m \in \{0, 1\}^\mu$, choose a randomness $r \xleftarrow{\$} \{0, 1\}^\tau$ and output $\text{Enc}_k(m; r) = \pi_k(m || r)$;
3. (decryption) Dec_k : on cipher $c \in \{0, 1\}^{\mu+\tau}$, output the first μ bits of the string produced by $\text{Dec}_k(c) = \pi_k^{-1}(c) = m || r$;
4. (correctness) $(\text{Dec}_k \circ \text{Enc}_k)(m; r) = \pi_k^{-1}(\pi_k(m || r))_\mu = m$.

Let us now prove the security under a non-adaptive quantum chosen-ciphertext attack. Again, the intuition is that, once the strong pseudorandom permutation is taken to be quantum-secure, the pre-challenge phase reveals at most a polynomial amount of evaluations of the pseudorandom permutation, despite the presence of a decryption oracle for the permutation and its inverse. Therefore, no adaptive access to a quantum encryption oracle is sufficient to succeed at the challenge.

Theorem 9.12. *If $\text{QPRP}_{m+\tau}$ is a family of strong pseudorandom permutations, then the scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ from Construction 9.11 is DeclIND-QCCA1-secure.*

Proof. As in the previous construction, our goal is to show that any QPT adversary \mathcal{A} wins the DeclND-QCCA1 security game with at most negligible probability. We introduce a sequence of indistinguishable hybrid games until we arrive at a security game in which the challenge is perfectly hidden and the adversary cannot win. First, we replace the strong QPRP from the original security game with a perfectly random permutation operating on bit strings of length $\mu + \tau$. According to QPRP security, we only negligibly affect the overall success probability. Next, we introduce a hybrid in which the challenger keeps track of all randomness values that are being used in the game and chooses a fresh value for each encryption. Finally, we proceed with a hybrid game in which we exploit the blindness of quantum algorithm towards relabeling at a single location Lemma 8.1. To this end, we relabel the same random permutation at a single location during the challenge phase in a way that is indistinguishable to the adversary, thus making it impossible to succeed for the remainder of the game. Consider the following sequence of indistinguishable hybrid games:

GAME 0: Original security game with QPRP π_k and challenge randomness r^* .

GAME 1: Replace the QPRP π_k with a random permutation $\pi \in S_{m+\tau}$.

GAME 2: Guarantee fresh randomness for each new encryption, else abort.

GAME 3: Relabel at the challenge phase with $s \xleftarrow{\$} \{0,1\}^{\mu+\tau}$:

$$\pi_s(x||y) = \begin{cases} s, & \text{for } y = r^* \\ \pi(x||y), & \text{for } y \neq r^*. \end{cases} \quad (9.13)$$

Suppose there exists a QPT adversary \mathcal{A} that wins GAME 0 with nonnegligible probability, hence there exists some polynomial $p(n)$ such that:

$$\Pr[\mathcal{A} \text{ outputs } b' = b] \geq 1/2 + 1/p(n). \quad (9.14)$$

Our claim is that the same adversary must (up to at most negligible probability) also succeed at GAME 3 by means of a hybrid argument.

GAME 0 vs. GAME 1:

Since we assumed the strong QPRP π_k to be quantum-secure, we can replace it with a random permutation π on the bit strings of length $\mu + \tau$ and only negligibly affect the success probability in Eq.(9.14). Note that an adversary \mathcal{A} that succeeds at GAME 0 but not at GAME 1 (i.e. \mathcal{A} only wins with at most negligible probability), allows us to build a π_k -distinguisher that violates QPRP security when receiving quantum oracle access to \mathcal{O}_φ and $\mathcal{O}_{\varphi^{-1}}$, where φ is a function $\varphi : \{0,1\}^{\mu+\tau} \rightarrow \{0,1\}^{\mu+\tau}$. As in the previous proof, we can construct a distinguisher \mathcal{D} by running \mathcal{A} such that and for infinitely many n :

$$\left| \Pr_{k \xleftarrow{\$} \{0,1\}^n} [\mathcal{D}^{|\pi_k\rangle|\pi_k^{-1}\rangle}(1^n) = 1] - \Pr_{\pi \xleftarrow{\$} S_{m+\tau}} [\mathcal{D}^{|\pi\rangle|\pi^{-1}\rangle}(1^n) = 1] \right| \geq 1/p(n) - \epsilon(n). \quad (9.15)$$

Therefore, our crucial assumption of strong QPRP security is broken by this distinguisher and we conclude that \mathcal{A} must also win GAME 1 with nonnegligible probability.

GAME 1 vs. GAME 2:

Since the challenger keeps track of all randomness values \mathcal{R} previously used for encryption, the adversary is guaranteed to receive fresh randomness in $\{0, 1\}^n \setminus \mathcal{R}$, else the game is aborted. Moreover, as the probability of sampling the same randomness twice is negligible, the difference between these two games is only negligible with regard to the overall success probability in the security game.

GAME 2 vs. GAME 3:

Again, as in the proof for the previous construction, the challenge from GAME 2 looks uniformly random. However, since the adversary has quantum oracle access to the random permutation π , it is possible to generate superpositions over the entire input space. Here, these states contain quantum information on the challenge value $\pi(\cdot || r^*)$ and our goal is to show that this advantage is still negligible. We make use of a blinding argument and argue that for any quantum adversary playing DecIND-QCCA1 , GAME 2 and GAME 3 are indistinguishable. Thus, up to negligible probability, the view of the adversary is that of GAME 3 in which the challenge is perfectly hidden and he cannot succeed.

Any quantum adversary playing the DecIND-QCCA1 game against Construction 9.11 can be separated into a pre-challenge routine $\mathcal{A}_0(1^n)$, as well as final challenge routine $\mathcal{A}_1(1^n)$. Upon an initial state $|\Psi_0\rangle$, $\mathcal{A}_0(1^n)$ performs Q unitary computations U_1, U_2, \dots, U_Q and alternating queries to the encryption and decryption oracle and outputs:

$$|\Psi^\pi\rangle = U_Q \mathcal{O}_{Dec} U_{Q-1} \mathcal{O}_{Enc} U_{Q-2} \dots U_2 \mathcal{O}_{Dec} U_1 \mathcal{O}_{Enc} U_0 |\Psi_0\rangle. \quad (9.16)$$

After the challenge, $\mathcal{A}_1(1^n)$ applies final encryption oracle queries with respect to π , as well as final unitary computations $U_{Q+1}, U_{Q+2}, \dots, U_T$ upon $|\Psi^\pi\rangle$ and generates:

$$|\Psi_\pi^\pi\rangle = U_T \mathcal{O}_{Enc} U_{T-1} \dots U_{Q+2} \mathcal{O}_{Enc} U_{Q+1} |\Psi^\pi\rangle, \quad (9.17)$$

where T is the total number of queries. Finally, after completing both routines, \mathcal{A} measures the final output state and outputs b' .

Suppose now that the adversary \mathcal{A} succeeds at GAME 2 but not at GAME 3, i.e. \mathcal{A} only wins with at most negligible probability. We argue that the same adversary must then win the RelabelingGame , hence must violate Proposition 8.3.

Let \mathcal{D} be the distinguisher playing the RelabelingGame given quantum oracle access \mathcal{O}_φ to a function $\varphi : \{0, 1\}^{\mu+\tau} \rightarrow \{0, 1\}^{\mu+\tau}$ whose goal is to output 0, whenever φ is a perfectly random function and to output 1, if $\varphi = \pi_s$ for some $x^* \xleftarrow{\$} \{0, 1\}^\mu$ and $s \xleftarrow{\$} \{0, 1\}^{\mu+\tau}$. \mathcal{D} now proceeds as follows: It simulates the IND-QCCA1 security game of \mathcal{A} by responding to quantum queries using oracles Enc_φ and Dec_φ that can be implemented based on the output of \mathcal{O}_φ . If \mathcal{A} wins the DecIND-QCCA1 and outputs $b' = b$, then \mathcal{D} outputs 1 and outputs 0 else. Note that, if $\varphi = \pi$, then the output of \mathcal{D} when running \mathcal{A} is identical to that of \mathcal{A} in GAME 1. Similarly, if $\varphi = \pi_s$, then the output is identical to that of \mathcal{A} in

GAME 2. The adversary \mathcal{A} is efficient, hence \mathcal{D} is also efficient by construction. Moreover, since we assumed that \mathcal{A} succeeds with nonnegligible probability in GAME 2, so does the distinguisher \mathcal{D} , and for infinitely many n :

$$\left| \Pr_{s,r^*}[\mathcal{D}^{\lvert \pi_s \rangle}(1^n) = 1] - \Pr_{s,r^*}[\mathcal{D}^{\lvert \pi \rangle}(1^n) = 1] \right| \geq 1/p(n) - \epsilon(n). \quad (9.18)$$

Therefore, we conclude that \mathcal{A} must also win GAME 3 with nonnegligible probability, in violation of Proposition 8.3.

Finally, note that in GAME 3, as in the previous proof, the `DecIND` challenge phase again amounts to the impossible task of distinguishing between two uniformly random strings. Therefore, since the view of \mathcal{A} is indistinguishable from such a challenge of perfect secrecy, we must now arrive at an overall contradiction and conclude that no such `QPT` algorithm \mathcal{A} with nonnegligible success probability exists. In summary, we must therefore conclude that for any quantum algorithm \mathcal{A} playing the `DecIND-QCCA1` game:

$$\Pr[\mathcal{A} \text{ wins } \text{DecINDGame}] = \frac{1}{2} + \epsilon(n). \quad (9.19)$$

□

Consequently, again as a direct consequence of the equivalence results of Section 9.1.4, we can conclude the previous result with the following additional observation:

Corollary 9.13. *Let `QPRF` be a function family of quantum-secure pseudorandom functions. Then, $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ from Construction 9.11 is both `IND-QCCA1`-secure, as well as `SEM-QCCA1`-secure.*

10 The Physical Realization of Quantum Computation

The simulation of quantum systems turns out to scale surprisingly poorly on conventional classical computers. In order to simulate N spin 1/2 particles and to solve the Schrödinger equation, one needs to store vectors of size 2^N , as well as manipulate matrices of size $2^N \times 2^N$. Due to this exponential scaling, it is well known that classical computers are highly inefficient in simulating the dynamics of quantum systems. This fact has already puzzled physicists in the 1980s, who hypothesized that an intrinsically quantum mechanical computer could potentially be more suitable for these tasks. In fact, it is often attributed to Richard Feynman [Fey82] to have been the first to speculate on the possibility of building quantum computers. The prospect of building a computer that would outperform any classical computing architecture and efficiently simulate quantum physics seemed captivating. In 1989, David Deutsch gave the first example of a quantum algorithm for a black box problem which could be solved faster with quantum mechanical means than with classical ones [Deu89]. Perhaps most notably, it was Peter Shor's 1994 discovery of efficient quantum algorithms for the factoring of integers and computing discrete logarithms [Sho94] that truly drew the attention towards the field of quantum computation. Only a few years later, Ignacio Cirac and Peter Zoller proposed a physical system of trapped ions on which quantum information processing could be realized [CZ95]. In this architecture, single trapped ions are engineered to carry quantum information and are both manipulated and measured with focused laser beams. Already within a year's time, David Wineland's group at National Institute of Standards and Technology achieved a breakthrough in ion-trap quantum computers [MMKW99], namely a controlled bit flip on a single ion. This experiment is often considered as the birth of experimental quantum computation. In this chapter, we give a basic introduction to trapped-ion quantum computers. To this end, we follow an excellent survey on trapped-ion computation by Häffner and Blatt [HRB08]. In later sections, we also describe recent implementations of quantum algorithms and further advances in the field.

10.1 DiVincenzo Criteria

All quantum information processing is concerned with the storage and coherent manipulation of information in a quantum system. In the previous chapters, we showed how quantum computers could solve certain mathematical problems faster than classical computers using the principles of quantum mechanics. In order to harness this quantum speed-up, however, one has to realize quantum computation in a physical system. Fortunately, nature presents us with many possible ways of realizing a qubit in a physical system. As typical representations of a qubit are found in the two states of a spin 1/2 particle, the vertical or horizontal polarization of a photon or simply the ground and excited states of an atom, each representation comes with its own drawbacks and advantages. While photons are easy to generate, they have proven to be difficult to interact on the basis of nonlinear materials alone [NC10]. Similarly, both the observation and control of spin states poses great difficulty, unless a carefully engineered environment is achieved. An example of such

circumstances is realized in a trapped-ion quantum computer, where ions are confined in a potential trap and subsequently cooled.

In 1996, David DiVincenzo, at the IBM Thomas J. Watson Research Center, proposed the following list of guidelines for a successful physical implementation of a quantum computer: [DiV00]

1. A scalable physical system with well characterized qubits.
2. The ability to initialize the state of the qubits to a simple initial state.
3. Long relevant decoherence times, much longer than the gate operation time.
4. A universal set of quantum gates.
5. A qubit-specific measurement capability.

Having the possibility of a functioning interface between quantum computers and devices for quantum communication in mind, DiVincenzo also added two additional requirements:

6. The ability to interconvert stationary and flying qubits.
7. The ability to faithfully transmit flying qubits between specified locations.

Currently, the trapped-ion computer is oftentimes regarded as the leading quantum computing architecture, while the runner-up technology is believed to be that of the solid-state architecture of superconducting qubits [LMRD17]. In this thesis, we present the ion-trap quantum computer as a model for quantum computation and discuss its fundamental properties and capabilities. When performing quantum information processing, such as the algorithms from the previous chapters, the ability to coherently manipulate as well as store information with low rates of error is crucial. Decoherence of quantum systems poses enormous difficulty to both of these tasks. In the next section, we will outline the extent to which trapped ion computation satisfies these criteria.

10.2 Ion-Trap Implementation

In this section, we discuss the physical realization of quantum computation on the basis of the ion-trap, the most successful quantum computing architecture to date. As the representation of a qubit is found in the hyperfine levels of an ion, we begin with a section on the hyperfine structure of atoms and continue with the experimental setup in the subsequent chapter. In the following chapters thereafter, we discuss how the ion-trap quantum computer satisfies DiVincenzo's Criteria, as well as the extent to which all the necessary ingredients for the implementation of quantum algorithms of the previous chapters are realized. In particular, we show how to perform elementary single-qubit and two-qubit gates using focused laser beams. Finally, in the last section, we present a recent performance comparison between a state-of-the-art solid-state device running the same algorithms.

10.2.1 Hyperfine Structure

In order to realize a qubit as a physical carrier of information, one has to represent it in an appropriate two-level quantum system. Following DiVincenzo's criteria, the task is to define a qubit that is not only well-characterized, but can also be controlled and manipulated for the purpose of information processing. In the ion-trap quantum computer, a qubit is found in the internal atomic states of the ion. Although a single trapped ion features a broad energy landscape, sophisticated use of lasers allows us to isolate just two levels in the energy spectrum of the atom. Alkali atoms present a popular choice for ion-trap experiments, as they feature a single valence electron in the outer shell, thus offering a simple and well-studied electronic structure. Typical ion candidates are the alkaline earth metals ${}^9\text{Be}^+$, ${}^{24}\text{Mg}^+$, ${}^{40}\text{Ca}^+$, as well as ${}^{171}\text{Yb}^+$, which, once ionized, behaves quite similarly. Each ion comes with a different mass and electronic transition at a certain wavelength, both highly relevant factors for trapping, as well as laser manipulation. For example, while lighter ions carry less inertia and are therefore easier to trap, they tend to exhibit electronic transitions at wavelengths in the deep ultra-violet that are less suitable for fiber-optics.

At high resolution, the atomic spectrum is known to feature a splitting of energy levels into further substructures, the so-called *fine structure* and *hyperfine structure*. Neither of the two structures are explained in the original Bohr model or predicted by Schrödinger theory and result from spin contributions of the electron spin and nuclear spin. The atomic states relevant for the representation of a qubit result from the sum of electron spin S and nuclear spin I , giving a total of $F = S + I$, where F is the total angular momentum. Using the long-lived states of the hyperfine structure, especially long coherence times can be achieved which, in this regard, make ion traps an ideal choice for a quantum computer.

Particularly the ytterbium isotope ${}^{171}\text{Yb}^+$ has become a favorable choice in recent experiments [DLFL16][FHM16] due to its large hyperfine splitting and strong ${}^2\text{S}_{1/2} \leftrightarrow {}^2\text{P}_{1/2}$ electronic transition around a wavelength of 369.53nm. An example of the hyperfine structure of ${}^{171}\text{Yb}^+$ is shown in Figure 15. This particular isotope exhibits long trapping lifetimes and, due to its strong electronic transition, it is well suited for broadband laser manipulation, as well as integration with optical fibers.

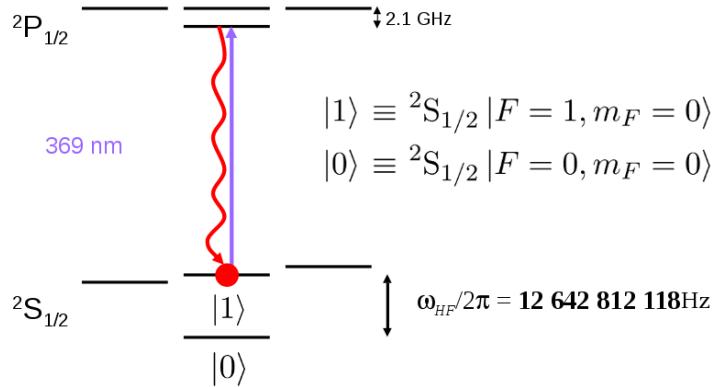


Figure 15: ([Mon13]) The spin contribution to the atomic energy levels of $^{171}\text{Yb}^+$. The hyperfine splitting results from the interaction between the spin-1/2 electron and the spin-1/2 nucleus.

Typically, the qubit is taken to be the two (first-order magnetic field-insensitive) hyperfine levels of the $^2\text{S}_{1/2}$ ground state [OYM07]:

$$\begin{aligned} |0\rangle &\equiv {}^2\text{S}_{1/2} |F = 0, m_F = 0\rangle \\ |1\rangle &\equiv {}^2\text{S}_{1/2} |F = 1, m_F = 0\rangle. \end{aligned}$$

Here, F and m_F denote the quantum numbers associated with the total angular momentum and its projection along the quantization axis defined by an applied magnetic field of 5.2 G . Since the magnetic quantum number is $m_F = 0$, the two hyperfine states carry only a quadratic *Zeeman shift*. Consequently, the $^{171}\text{Yb}^+$ ion features a particular insensitivity with respect to magnetic field fluctuations. Often in the ion-trap literature, the notation $|g\rangle = |0\rangle$ and $|e\rangle = |1\rangle$, denoting the ground and excited states respectively, is adopted in order to avoid confusion around additional coupling with vibrational modes of the ion chain. The qubit frequency splitting between the above $^2\text{S}_{1/2}$ states is in the order of $\nu_0 = 12.642821\text{ GHz}$. Most notably, Monroe et al. [OYM07] have measured average qubit coherence times of $2.5(3)\text{s}$ that are significantly longer than the typical gate operation time at microseconds.

10.2.2 Experimental Setup

The main component of an ion-trap quantum computer is an electromagnetic trap, a vacuum chamber surrounded by four cylindrical electrodes. In order to produce the necessary trapping potential towards axial confinement of the ions, the end caps of the rods are biased at different voltages. However, as we will show now, trapping ions by means of static electric fields alone is not possible. *Earnshaw's theorem* states that a charged particle cannot be confined in three dimensions by static electric fields, as the divergence of the field vanishes in empty space. This fact can easily be verified in the following short argument. Let $\mathbf{r}_0 = (x_0, y_0, z_0)$ denote the coordinates of the charge. Then, for any trapping potential, we require that the particle returns to its equilibrium position once it is displaced.

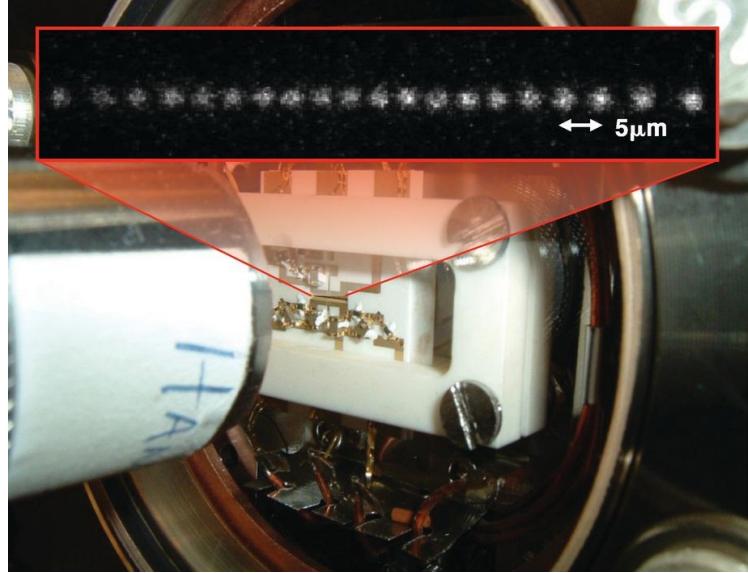


Figure 16: ([Mon13]) The University of Maryland ion-trap at the Chris Monroe lab, 2013. Each dot represents a single $^{171}\text{Yb}^+$ ion exhibiting state dependent fluorescence when driven by individual focused laser beams. A CCD camera collects fluorescence from the scattering of photons and creates an image over thousands of measurements.

Consequently, we demand from the potential energy $U(\mathbf{r})$ that:

$$\nabla U(\mathbf{r}_0) = 0, \quad \nabla^2 U(\mathbf{r}_0) > 0. \quad (10.1)$$

However, since the electric potential energy is given by $U(\mathbf{r}) = q\Phi(\mathbf{r})$, where Φ is the electrostatic potential, we must conclude from *Gauss' law* that:

$$\nabla^2 U(\mathbf{r}) = q \nabla^2 \Phi(\mathbf{r}) = 0. \quad (10.2)$$

Thus, the electric potential obeys the *Laplace equation* and violates the previous condition Eq.(10.1) we required for the desired potential trap. In order to achieve confinement, one has to adopt time-varying electric fields that, on average, create an effective trap in three dimensions. In practice, this is can be realized in a *rf Paul trap* [POF58], where a combination of both static as well as oscillating electric fields switching at rates around radio frequency produce an effective harmonic trap. As a result, a potential of *quadropole geometry* is generated that confines a charge in all three dimensions. A linear rf Paul trap can also successfully confine several ions simultaneously along its trap axis (typically taken to be the \hat{z} -axis)[Pau90]. Since the vibrations of trapped ions around their equilibrium positions are strongly coupled due to the Coulomb interaction, motion of any one single ion induces a joint oscillation in all other ions. The Hamiltonian describing the motion of N confined ions together with the Coulomb repulsion is given by:

$$\mathcal{H} = \sum_{j=1}^N \frac{M}{2} \left(\frac{\hat{p}_j^2}{M^2} + \omega_x^2 x_j^2 + \omega_y^2 y_j^2 + \omega_z^2 z_j^2 \right) + \sum_{j=1}^N \sum_{i>j} \frac{e^2}{4\pi\epsilon_0 |\hat{r}_j - \hat{r}_i|}, \quad (10.3)$$

where M is the mass of a single ion and ω_x, ω_y and ω_z describe the frequency of oscillation along the respective directions. For the sake of simplicity, one typically considers a linear Paul trap design in which only a single motional direction along the trap axis is selected for in which all ions lie along the \hat{z} -axis. If the displacement due to the oscillations is much smaller than the spatial separation between the ions, we can describe the vibrations (i.e. phonons) in an harmonic oscillator approximation [WMI98]. At low temperature, the linear chain of ions freezes into a crystal where, for a quantum of vibrational energy $\hbar\omega_z$, the desired cooling requires both that $k_B T \ll \hbar\omega_z$, as well as that the thermal energy T drops below the energy difference of the two atomic levels. A chain of N ions exhibits various normal modes of vibration, both radial and axial, each at frequencies independent of N . The axial mode of lowest frequency is given by the *center-of-mass mode* (COM), a collective motion of the entire ion chain along the trap axis. In order to control and encourage such joint motion in the COM mode, it is necessary to suppress vibrational modes of higher frequency (such as relative or radial motion) by applying Doppler-cooling and preparing the ions in their motional ground state [WI79].

The use of resonant laser light is a fundamental component of the ion-trap computer and appears throughout multiple stages of quantum information processing, such as groundstate-cooling, qubit initialization, qubit gate-operations and state detection. In order to achieve state initialization, sophisticated use of *optical pumping* can drive hyperfine transitions into short-lived and energetically distant states that subsequently decay back to the ground state according to known *selection rules* (Figure 17). Measurement, or state detection, works using state dependent fluorescence as follows: If the qubit state is in the excited state $|1\rangle$, the 369.53 nm light applied for detection is nearly on resonance, and the ion exhibits fluorescence by scattering many photons. If, however, the state is in the ground state $|0\rangle$, very few photons are scattered and we observe a dark state. Finally, a photon count results in accurate state detection. Moreover, as we discuss in the subsequent chapters, manipulation of qubits can be realized as an optical Rabi oscillation under resonant laser light.

Let us conclude this section by briefly summarizing how ion-trap quantum computers fulfill the DiVincenzo criteria:

1. **A scalable physical system with well characterized qubits:** The atomic hyperfine states are exceptionally long lived and serve as an ideal choice for qubits, see [OYM07]. Though scalable in principle, complications typically arise in both mass and mode structure of sufficiently long ion chains. Modern techniques circumvent these problems and address scalability by means of ion-transport among multiple ion traps [WMI98][FHM16].
2. **The ability to initialize the state of the qubits to a simple initial state:** State initialization is achieved using optical pumping, a technique that prepares hyperfine ground states with average fidelities > 0.99 , for example [OYM07].

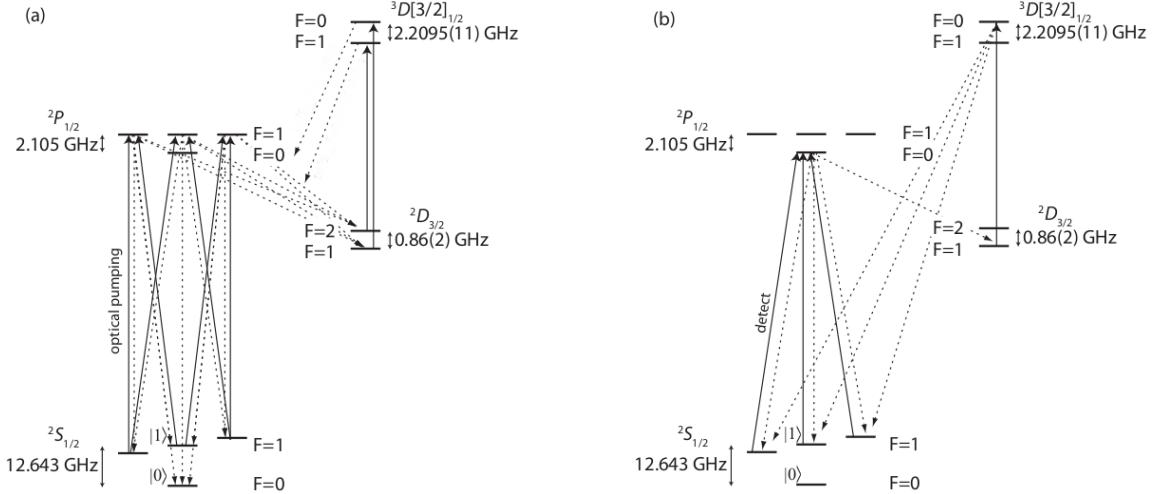


Figure 17: ([OYM07]) Optical pumping for state initialization (a) and state detection (b) of $^{171}\text{Yb}^+$. The nuclear spin is given by $I = 1/2$. Appropriate polarization of the incoming laser beam can exploit atomic selection rules and initialize the desired groundstate.

3. **Long relevant decoherence times, much longer than the gate operation time:** Typical coherence times of modern ion-trap architectures average around a few milliseconds and are therefore several orders of magnitude longer than the time scale required for quantum gate operations, see [OYM07].
4. **A universal set of quantum gates:** All single-qubit gates can be performed using laser pulses that drive Rabi oscillations between the two atomic levels. Two-qubit gates are implemented by exploiting the long range Coulomb interaction, such as in the original proposal by Cirac and Zoller [CZ95]. Quantum information from a single ion can be transferred into the common motional degree of freedom of the entire ion string using a sideband transition by focused laser pulses. Such conditional quantum dynamics are sufficient to give rise to elementary two-qubit gates needed for universal computation. Moreover, sources of error during larger scale quantum operations can be controlled for by more sophisticated types of multiparticle entanglement, such as Mølmer-Sørensen interactions [MS99].
5. **A qubit-specific measurement capability:** Measurements are performed using state dependent fluorescence in which photon scattering allows for state read-out of individual qubits.

10.2.3 The Hamiltonian

In this section, we discuss the basic Hamiltonian of a single trapped ion interacting with near resonant laser light. The two-level approximation is valid in this regime, as all other atomic levels are energetically far away and highly detuned. Similar to a spin-1/2 system under a time-dependent magnetic field, the two-level atom undergoes an optical Rabi oscillation under the action of the electromagnetic field.

Let us consider a Hamiltonian of a two-level system interacting with a quantized harmonic oscillator of vibrational modes through a laser beam, where:

$$\mathcal{H} = \mathcal{H}_{\text{atom}} + \mathcal{H}_{\text{free}} + \mathcal{H}_{\text{int}}. \quad (10.4)$$

Recall from the previous section that the Hamiltonian describing the free motion of a single ion along the trap axis in an effective harmonic potential can be written as:

$$\mathcal{H}_{\text{free}} = \frac{p^2}{2M} + \frac{1}{2}M\omega_z^2 z^2, \quad (10.5)$$

where ω_z is the frequency of oscillation around the equilibrium position in the \hat{z} direction. If the coupling to the external field is small and the ion inside the vacuum chamber is well isolated from its surroundings, its motion becomes quantized and we can introduce raising and lowering operators $z = \sqrt{\frac{\hbar}{2M\omega_z}}(a + a^\dagger)$ and $p = i\sqrt{\frac{\hbar M\omega_z}{2}}(a - a^\dagger)$. Consequently, together with the Hamiltonian corresponding to the internal atomic levels, we can write:

$$\mathcal{H}_{\text{atom}} = \hbar\omega_{eg} \frac{\sigma_z}{2} \quad (10.6)$$

$$\mathcal{H}_{\text{free}} = \hbar\omega_z \left(a^\dagger a + \frac{1}{2} \right). \quad (10.7)$$

In the following, we will denote $\mathcal{H}_0 = \mathcal{H}_{\text{atom}} + \mathcal{H}_{\text{free}}$. The Hamiltonian \mathcal{H}_{int} describes the atom-light interaction of the ion with the laser. Following Wineland et al. [WBB03], the interaction between the ion and the electric field of the laser beam is given by:

$$\mathcal{H}_{\text{int}}(t) = -\vec{d} \cdot \vec{E} = -\vec{d} \cdot E_0 \hat{\epsilon}_L \cos(kz - \omega_L t + \phi), \quad (10.8)$$

where \vec{d} is the electric dipole operator, \vec{E} is the (classical) electric field, E_0 the field strength, z is the position operator of the ion for displacement from its equilibrium position, $\hat{\epsilon}_L$ is the laser beam polarization, ω_L is the frequency of the laser, k is the laser beam's k -vector parallel to \hat{z} (the axis of the trap) and where ϕ is the phase of the laser at the mean position of the ion. In the dipole approximation, \vec{d} can be further expanded in terms of the internal states of the atom, since it is proportional to $\sigma_+ + \sigma_-$, where $\sigma_+ = |e\rangle\langle g|$ and $\sigma_- = |g\rangle\langle e|$. By introducing the Rabi flop frequency $\Omega = -\frac{E_0}{\hbar} \langle e | \vec{d} \cdot \hat{\epsilon}_L | g \rangle$ and the Lamb-Dicke parameter $\eta = k \sqrt{\frac{\hbar}{2M\omega_z}}$, we can express Eq.(10.8) as:

$$\mathcal{H}_{\text{int}}(t) = \hbar \frac{\Omega}{2} (\sigma_+ + \sigma_-) \left(e^{i(\eta(a+a^\dagger) - \omega_L t + \phi)} + e^{-i(\eta(a+a^\dagger) - \omega_L t + \phi)} \right). \quad (10.9)$$

Taking the width of the ion's oscillation along the trap axis at low temperatures to be small compared to the wavelength of the incoming laser beam, we can apply the Lamb-Dicke limit ($\eta \sqrt{\langle (a+a^\dagger)^2 \rangle} \ll 1$) and further expand the relevant exponential from Eq.(10.9):

$$e^{i\eta(a+a^\dagger)} = 1 + i\eta(a + a^\dagger) + \mathcal{O}(\eta^2). \quad (10.10)$$

It is now convenient to work in the interaction picture $\mathcal{H}'_{\text{int}} = e^{i\mathcal{H}_0 t/\hbar} \mathcal{H}_{\text{int}} e^{-i\mathcal{H}_0 t/\hbar}$. Using the *Baker-Campbell-Hausdorff lemma*,

$$e^{\alpha A} B e^{-\alpha A} = B + \alpha [A, B] + \frac{\alpha^2}{2!} [A, [A, B]] + \frac{\alpha^3}{3!} [A, [A, [A, B]]] + \dots, \quad (10.11)$$

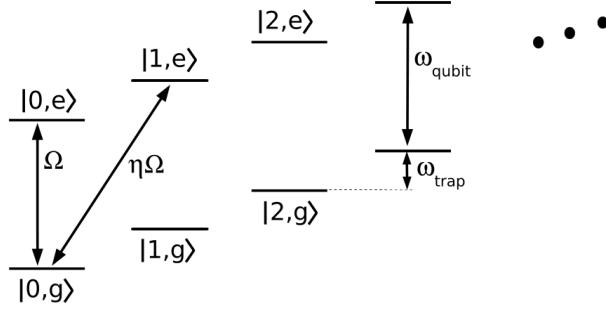


Figure 18: ([HRB08]) Transitions between atomic levels and phonon modes.

we get the following identities:

$$e^{i\omega_z a^\dagger a t} \left[1 + i\eta (a + a^\dagger) \right] e^{-i\omega_z a^\dagger a t} = 1 + i\eta (ae^{-i\omega_z t} + a^\dagger e^{i\omega_z t}) \quad (10.12)$$

$$e^{i\omega_{eg} \sigma_z t/2} \sigma_+ e^{-i\omega_{eg} \sigma_z t/2} = \sigma_+ e^{i\omega_{eg} t}. \quad (10.13)$$

By applying a rotating wave approximation and assuming near resonance $\Delta = \omega_L - \omega_{eg} \approx 0$, we ignore all rapidly oscillating terms of the form $\exp(\pm i(\omega_L + \omega_{eg})t)$ and find:

$$\mathcal{H}'_{\text{int}}(t) = \hbar \frac{\Omega}{2} \sigma_+ e^{-i(\Delta t - \phi)} \left[1 + i\eta (ae^{-i\omega_z t} + a^\dagger e^{i\omega_z t}) \right] + \text{h.c.} + \mathcal{O}(\eta^2). \quad (10.14)$$

A second rotating wave approximation assumes that only one transition at a time is considered, now gives the Hamiltonian:

$$\mathcal{H}'_{\text{int}}(t) = \frac{\hbar\Omega}{2} \left[\sigma_+ e^{-i(\Delta t - \phi)} + \sigma_- e^{i(\Delta t - \phi)} + i\eta (\sigma_+ e^{-i(\Delta t - \phi)} - \sigma_- e^{-i(\Delta t - \phi)}) (ae^{-i\omega_z t} + a^\dagger e^{i\omega_z t}) \right] \quad (10.15)$$

We can now identify three different cases of interest with respect to the detuning of the laser beam, the so-called carrier and sideband transitions:[HRB08]

1. The carrier transition ($\omega_L = \omega_{eg}$, $\Delta = 0$):

$$\mathcal{H}_C = \hbar \frac{\Omega}{2} (\sigma_+ e^{i\phi} + \sigma_- e^{-i\phi}). \quad (10.16)$$

In this regime, transitions $|g\rangle |n\rangle \leftrightarrow |e\rangle |n\rangle$ between the atomic states of the ion can be performed.

2. The blue sideband transition ($\omega_L = \omega_{eg} + \omega_z$, $\Delta = \omega_z$):

$$\mathcal{H}_+ = i\hbar \frac{\Omega}{2} \eta (\sigma_+ a^\dagger e^{i\phi} - \sigma_- a e^{-i\phi}). \quad (10.17)$$

This allows for the creation of a phonon mode and simultaneous excitation of the atomic state: $|g\rangle |n\rangle \leftrightarrow |e\rangle |n+1\rangle$.

3. The red sideband transition ($\omega_L = \omega_{eg} - \omega_z$, $\Delta = -\omega_z$):

$$\mathcal{H}_- = i\hbar\frac{\Omega}{2}\eta(\sigma_+ae^{i\phi} + \sigma_-a^\dagger e^{-i\phi}). \quad (10.18)$$

Simultaneous to exciting the atomic state of the ion, a phonon mode is destroyed. Thus the following transitions can be performed: $|g\rangle|n\rangle \leftrightarrow |e\rangle|n-1\rangle$.

Note that the red sideband Hamiltonian is formally equivalent to the well-known *Jaynes-Cummings* Hamiltonian in quantum optics that describes a two-level atom interacting with a quantized mode of an optical cavity. Since the Coulomb interaction provides a strong coupling among the ions, the entire chain of ions exhibits various normal modes of motion, each at different frequencies, such as center-of-mass mode, the stretch mode or the axial mode [HRB08]. In order to describe the full Hamiltonian of a linear crystal consisting of N ions, we can introduce a sum over all single ion contributions and respective vibrational modes of the entire ion chain, as follows:

$$\mathcal{H}_0 = \sum_{j=1}^N \hbar\omega_{eg} \frac{\sigma_{zj}}{2} + \sum_{l=1}^N \hbar\omega_{z_l} \left(a_l^\dagger a_l + \frac{1}{2} \right) \quad (10.19)$$

$$\mathcal{H}'_{\text{int}} = \sum_{j=1}^N \frac{\hbar\Omega_j}{2} \sigma_{+j} e^{-i(\Delta t - \phi)} \exp \left(i \sum_{l=1}^N \eta_{jl} [a_l e^{-i\omega_z t} + a_l^\dagger e^{i\omega_z t}] \right) + \text{h.c.} \quad (10.20)$$

Repeating the analysis of the single-ion Hamiltonian, we can write the interaction Hamiltonian in the rotating wave approximation and the Lamb-Dicke limit as:

$$\mathcal{H}'_{\text{int}} = \sum_{j,l=1}^N \frac{\hbar\Omega_j}{2} \left[\sigma_+^{(j)} e^{-i(\Delta t - \phi)} + \sigma_-^{(j)} e^{i(\Delta t - \phi)} + i\eta_{jl} (\sigma_+^{(j)} e^{-i(\Delta t - \phi)} - \sigma_-^{(j)} e^{i(\Delta t - \phi)}) (a_l e^{-i\omega_z t} + a_l^\dagger e^{i\omega_z t}) \right] \quad (10.21)$$

Here, we applied ground-state cooling and prepared only the lowest frequency COM mode in which the same phonon is shared among all ions in the crystal. In this regime, we can implement a two-qubit gate using the common motional degree of freedom as a *bus* to transfer conditional information among the ions. In the next sections, we describe how to realize single-qubit gates, as well as two-qubit gates, in the ion-trap quantum computer.

10.2.4 Single-Qubit Gates

In Chapter 4.4, we discussed how all quantum operations can be broken down into a sequence of single qubit and two-qubit operations. A major advantage of trapped-ion quantum computers lies in the fact that single-qubit operations are particularly easy to implement, as well as to control through the use of resonant laser light. In fact, we can show that any single-qubit operation corresponds to a rotation on the Bloch sphere and can thus be realized as a Rabi oscillation between the two qubit levels using a resonant laser pulse. In practice, such tuning of appropriate pulse parameters takes place at the control interface given by an acousto-optical modulator (AOM) [DHL05].

Consider a two-level system that starts out in some internal state $|\psi\rangle = c_g|g\rangle + c_e|e\rangle$ at time $t = 0$. Under a stationary Hamiltonian, the subsequent time-evolution after time τ is governed by the unitary dynamics:

$$|\psi(\tau)\rangle = \exp\left(\frac{-i\mathcal{H}\tau}{\hbar}\right)|\psi(0)\rangle. \quad (10.22)$$

Considering the stationary Hamiltonians \mathcal{H}_C , \mathcal{H}_+ and \mathcal{H}_- from the previous section under radiation of pulse length τ and respective detuning, we arrive at unitary dynamics which induce the following rotations:

$$R_C(\theta, \phi) = \exp\left(-i\theta/2(\sigma_+e^{i\phi} + \sigma_-e^{-i\phi})\right) \quad (10.23)$$

$$R_+(\theta, \phi) = \exp\left(-i\theta/2(\sigma_+a^\dagger e^{i\phi} - \sigma_-a e^{-i\phi})\right) \quad (10.24)$$

$$R_-(\theta, \phi) = \exp\left(-i\theta/2(\sigma_+a e^{i\phi} + \sigma_-a^\dagger e^{-i\phi})\right), \quad (10.25)$$

where the control parameters $\theta = \Omega\tau$ (or $\theta = \Omega\eta\tau$ for the sideband evolution) and phase ϕ determine the nature of the rotation. Note that the phase parameter ϕ of the laser at the start of the interaction experiment is completely arbitrary but sets the reference for all subsequent operations. We can identify the result of any of the above dynamics by a rotation operator $R(\theta, \phi)$ acting on $|\psi\rangle$ in terms of a rotation in the equatorial plane by ϕ and a rotation θ in the vertical plane. For example, in the case of the carrier evolution, this allows us to decompose the evolution as:

$$R_C(\theta, \phi) = \exp\left(-i\theta/2(\sigma_+e^{i\phi} + \sigma_-e^{-i\phi})\right) \quad (10.26)$$

$$= \mathbb{1} \cos \theta/2 - i(\sigma_x \cos \phi - \sigma_y \sin \phi) \sin \theta/2 \quad (10.27)$$

$$= \begin{pmatrix} \cos \theta/2 & -ie^{i\phi} \sin \theta/2 \\ -ie^{-i\phi} \sin \theta/2 & \cos \theta/2 \end{pmatrix} \quad (10.28)$$

Thus, by fixing ϕ appropriately, we can now identify the following set of rotation operators in the x and y plane:

$$R_x(\theta) = \begin{pmatrix} \cos \theta/2 & -i \sin \theta/2 \\ -i \sin \theta/2 & \cos \theta/2 \end{pmatrix} \quad (10.29)$$

$$R_y(\theta) = \begin{pmatrix} \cos \theta/2 & -\sin \theta/2 \\ \sin \theta/2 & \cos \theta/2 \end{pmatrix} \quad (10.30)$$

In order to obtain $R_z(\theta)$, we can use a natural decomposition into rotations around the x and y axis by writing $R_z(\theta) = R_y(\frac{\pi}{2})R_x(\theta)R_y(-\frac{\pi}{2})$. Thus, rotations around the z axis are given by the rotation operator:

$$R_z(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix} \quad (10.31)$$

In fact, *any* unitary single-qubit operation U can be decomposed using the rotation operators above, as stated in Theorem 4.1. Consider, for example, a resonant pulse of length

$\Omega\tau = \pi$ which realizes a 180° rotation (up to an overall phase):

$$R_x(\pi) |\psi\rangle = -i\sigma_x |\psi\rangle. \quad (10.32)$$

Another important gate is the Hadamard gate, which we can now realize as a $\frac{\pi}{2}$ -pulse in the y -plane and write $H = R_y(\pi/2)$. For example, given the initial state $|g\rangle$, we can easily create an equal superposition by performing a Hadamard $\frac{\pi}{2}$ -pulse:

$$H |g\rangle = R_y(\pi/2) |g\rangle = \frac{|g\rangle + |e\rangle}{\sqrt{2}}. \quad (10.33)$$

Starting from an initial state $|g\rangle$, we can also prepare any pure state $|\psi\rangle$ on the Bloch sphere (Figure 1) by an appropriate choice of control parameters $\theta = \Omega\tau$ and ϕ using the unitary dynamics in Eq.(10.26):

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right) |g\rangle + e^{i\phi} \sin\left(\frac{\theta}{2}\right) |e\rangle. \quad (10.34)$$

If the laser is slightly detuned, we can repeat the analysis above for sideband rotations that at the same time increase the vibrational modes. Notice that now the control parameters are $\theta = \Omega\eta\tau$ and ϕ .

10.2.5 Two-Qubit Gates

According to early work by David Deutsch [Deu89], a universal set of gates can be achieved using single-qubit and two-qubit gates only. In the previous section, we introduced the means to generate single-qubit operations under laser radiation and subsequent Rabi oscillation. In order to describe a system of a linear chain of ions, each mutually coupled with the Coulomb interaction, one has to adopt a Hamiltonian that includes total contribution of all ions.

An early proposal for a two-qubit gate can already be found in the Cirac and Zoller [CZ95] design of the ion-trap computer. The idea of the Cirac-Zoller-gate is the following: A red sideband pulse onto the first ion transfers information from the atomic state into the motional degree of freedom, conditioned on its state. Once the ion begins oscillating, it affects the entire string of ions due to the strong Coulomb repulsion. Thus, the second ion can now be addressed with operations that are conditioned on the motional state of the first ion. Finally, another red sideband transition reverses the motional state and causes the first ion to return to its original state. The procedure works as follows:

1. A laser beam tuned to the red sideband frequency $\omega_{eg} - \omega_z$ and length $\theta = \pi$ is focused on the first ion. Depending on the atomic state of the ion, a transfer into the motional degree of freedom may occur. Consequently, if the ion starts out in the ground state, no state transfer occurs due to the detuning. If the ion is in an excited state, a phonon mode is created:

$$\begin{aligned} |g\rangle |0\rangle &\longrightarrow |g\rangle |0\rangle \\ |e\rangle |0\rangle &\longrightarrow -i |g\rangle |1\rangle. \end{aligned}$$

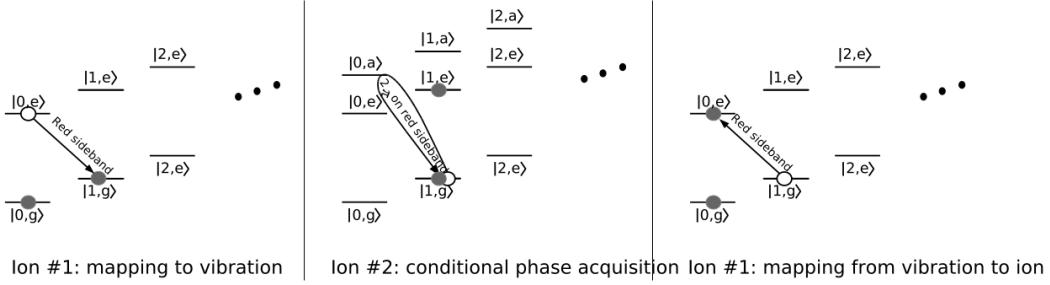


Figure 19: ([CZ01]) The two-qubit Cirac-Zoller gate.

2. A second laser tuned to the red sideband frequency with duration $\theta = 2\pi$ is now focused onto the second ion. This induces a 2π rotation between $|g\rangle|1\rangle$ and an auxiliary hyperfine state $|a\rangle|0\rangle$. Note that the design of the transition is such that all other states $|g\rangle|0\rangle$, $|e\rangle|0\rangle$, and $|e\rangle|1\rangle$ are left untouched, as there is insufficient energy to drive any of these levels. As a result, the following qubit operation is performed at the second ion:

$$\begin{aligned}
 |e\rangle|0\rangle &\longrightarrow |e\rangle|0\rangle \\
 |e\rangle|1\rangle &\longrightarrow |e\rangle|1\rangle \\
 |g\rangle|0\rangle &\longrightarrow |g\rangle|0\rangle \\
 |g\rangle|1\rangle &\longrightarrow -|g\rangle|1\rangle
 \end{aligned}$$

3. A final laser beam tuned to the red sideband frequency $\omega_{eg} - \omega_z$ and length $\theta = \pi$ is focused on the first ion to remove the motional quantum and restore the first ion to its original state.

$$\begin{aligned}
 |g\rangle|0\rangle &\longrightarrow |g\rangle|0\rangle \\
 |g\rangle|1\rangle &\longrightarrow -i|e\rangle|0\rangle.
 \end{aligned}$$

In summary, the Cirac-Zoller gate performs the following two-qubit operation:

$$\begin{array}{llll}
 R_-^{(1)}(\pi, 0) & R_-^{(2)}(2\pi, 0) & R_-^{(1)}(\pi, 0) \\
 \begin{array}{llll}
 |g\rangle|g\rangle|0\rangle & \longrightarrow & |g\rangle|g\rangle|0\rangle & \longrightarrow & |g\rangle|g\rangle|0\rangle \\
 |g\rangle|e\rangle|0\rangle & \longrightarrow & |g\rangle|e\rangle|0\rangle & \longrightarrow & |g\rangle|e\rangle|0\rangle \\
 |e\rangle|g\rangle|0\rangle & \longrightarrow & -i|g\rangle|g\rangle|1\rangle & \longrightarrow & |e\rangle|g\rangle|0\rangle \\
 |e\rangle|e\rangle|0\rangle & \longrightarrow & -i|g\rangle|e\rangle|1\rangle & \longrightarrow & -|e\rangle|e\rangle|0\rangle
 \end{array} & \begin{array}{llll}
 |g\rangle|g\rangle|0\rangle & \longrightarrow & |g\rangle|g\rangle|0\rangle & \longrightarrow & |g\rangle|g\rangle|0\rangle \\
 |g\rangle|e\rangle|0\rangle & \longrightarrow & |g\rangle|e\rangle|0\rangle & \longrightarrow & |g\rangle|e\rangle|0\rangle \\
 |e\rangle|g\rangle|1\rangle & \longrightarrow & i|g\rangle|g\rangle|1\rangle & \longrightarrow & |e\rangle|g\rangle|0\rangle \\
 |e\rangle|e\rangle|1\rangle & \longrightarrow & -i|g\rangle|e\rangle|1\rangle & \longrightarrow & -|e\rangle|e\rangle|0\rangle
 \end{array} & \begin{array}{llll}
 |g\rangle|g\rangle|0\rangle & \longrightarrow & |g\rangle|g\rangle|0\rangle & \longrightarrow & |g\rangle|g\rangle|0\rangle \\
 |g\rangle|e\rangle|0\rangle & \longrightarrow & |g\rangle|e\rangle|0\rangle & \longrightarrow & |g\rangle|e\rangle|0\rangle \\
 |e\rangle|g\rangle|0\rangle & \longrightarrow & |e\rangle|g\rangle|0\rangle & \longrightarrow & |e\rangle|g\rangle|0\rangle \\
 |e\rangle|e\rangle|0\rangle & \longrightarrow & -|e\rangle|e\rangle|0\rangle & \longrightarrow & -|e\rangle|e\rangle|0\rangle
 \end{array}
 \end{array}$$

The operation, as shown above, realizes the controlled-Z (CZ) gate. In the 2-qubit representation, it can be written as a unitary matrix:

$$U = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \quad (10.35)$$

By using a Ramsey-type experiment with two additional single-qubit $\pi/2$ pulses, the CZ-gate is easily turned into a CNOT gate, as follows:

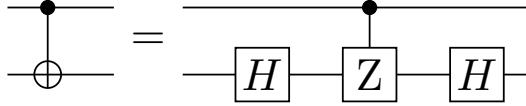


Figure 20: Using a CZ gate and two $\pi/2$ -pulses that perform Hadamard gates, one can construct a CNOT gate.

In the next section, we discuss recent implementations of many of the algorithms we present in this thesis.

10.2.6 Quantum Algorithms with Trapped Ions

Ever since the first quantum algorithms emerged after Deutsch's algorithm was first proposed, many algorithms have in fact been implemented on a quantum computer. In this respect, the ion-trap quantum computer still largely dominates all other architectures due to its long coherence times, a fact we discuss in the next section. Deutsch's algorithm was first successfully implemented as early as 2003, using $^{40}\text{Ca}^+$ ions [GHRL03]. Typical fidelities on identifying the function classes exceeded over 0.9.

The most comprehensive report on the implementation of quantum algorithms up to date was recently published by Monroe et al. at the University of Maryland [DLFL16]. Using a linear chain of five $^{171}\text{Yb}^+$ hyperfine qubits, a programmable interface allows the implementation of the Deutsch-Josza and Bernstein-Vazirani algorithm, Simon's algorithm and the quantum Fourier transform. It was highlighted in the publication that, compared to other architectures such as the solid-state implementations, the ion-trap quantum computer is more flexible since it is easy to program by external fields and can thus be reconfigured to run any of the above algorithms.

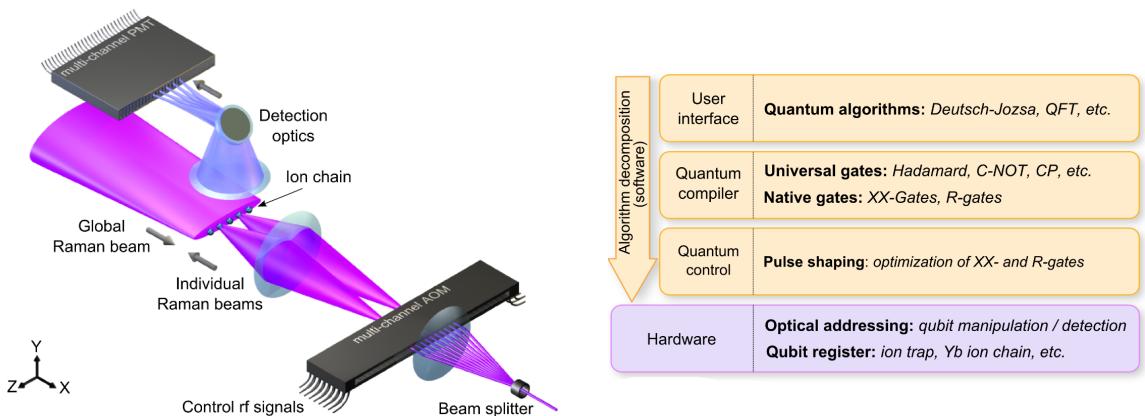


Figure 21: ([DLFL16]) The Maryland ion-trap setup. A user-interface is provided that allows versatile programming of a five-qubit ion-trap computer to run any algorithm.

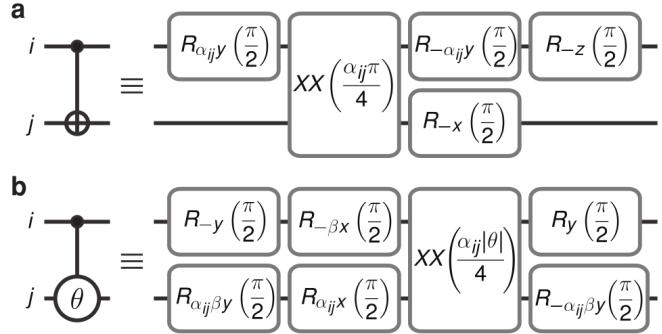


Figure 22: ([DLFL16]) Decomposition of two-qubit gates. Here (a) refers to the CNOT gate and (b) refers to the CP gate. In order to achieve two-qubit operations that are less prone to errors, Monroe et al. adopt decompositions into multiple single-qubit gates prior and after the two-qubit (XX)-gates based on Mølmer-Sørensen interactions [MS99].

The setup of the programmable ion-trap computer is as follows: At the top of the hierarchy, we find a flexible interface that allows a user to program the specifications of the desired algorithm. Here, a standard set of universal gates such as the Hadamard, the CNOT or the CP gate are available for programming. In analogy to a classical compiler, a quantum compiler translates these gates to a set of native gate instructions consisting of single-qubit rotation pulses or two-qubit Ising-like gates due to Mølmer-Sørensen interactions [MS99]. All native gates are finally performed as external light pulses originating from the acousto-optical modulator (AOM). In the usual setup, all ions are confined inside a linear rf-Paul trap and cooled near their motional ground state using Doppler cooling. As a result, the chain of ions freezes into a linear crystal, with equal spacing of around $5\mu\text{m}$. Next, the use of lasers and optical pumping achieves efficient state initialization, as described in [OYM07]. All quantum gate implementations follow coherent rotations using Raman transitions to drive both, atomic transitions, as well as vibrational transitions, in which lasers are focused on all of the ions in the chain simultaneously. Thus, in order to address ions individually, each Raman beam is split into a static array of beams processed at the AOM, which then focuses the beams onto the individual ions. Measurement is the result of driving transitions near 369nm of wavelength and simultaneously collecting state dependent fluorescence from each ion. In practice, this is done by using a multi-channel photo-multiplier tube (PMT). Thus, if the qubit is in the state $|1\rangle$, the laser is on resonance and state-dependent fluorescence can be collected. Else, if the qubit is in the state $|0\rangle$, the laser is sufficiently detuned and a dark state is observed.

At the bottom of the hierarchy, qubit operations are performed via pairs of Raman beams from a single 355nm YAG mode-locked laser. Here, the single qubit rotations $R_\phi(\theta)$ are performed by a Raman beat-note of defined amplitude, phase and duration at the qubit resonance frequency $\nu_0 = 12.642821\text{GHz}$. As introduced in Section 10.2.4, θ describes the rotation angle and ϕ is determined by the duration and phase-offset of the beat-note and is programmed at the appropriate AOM channels. The two-qubit gates are performed using

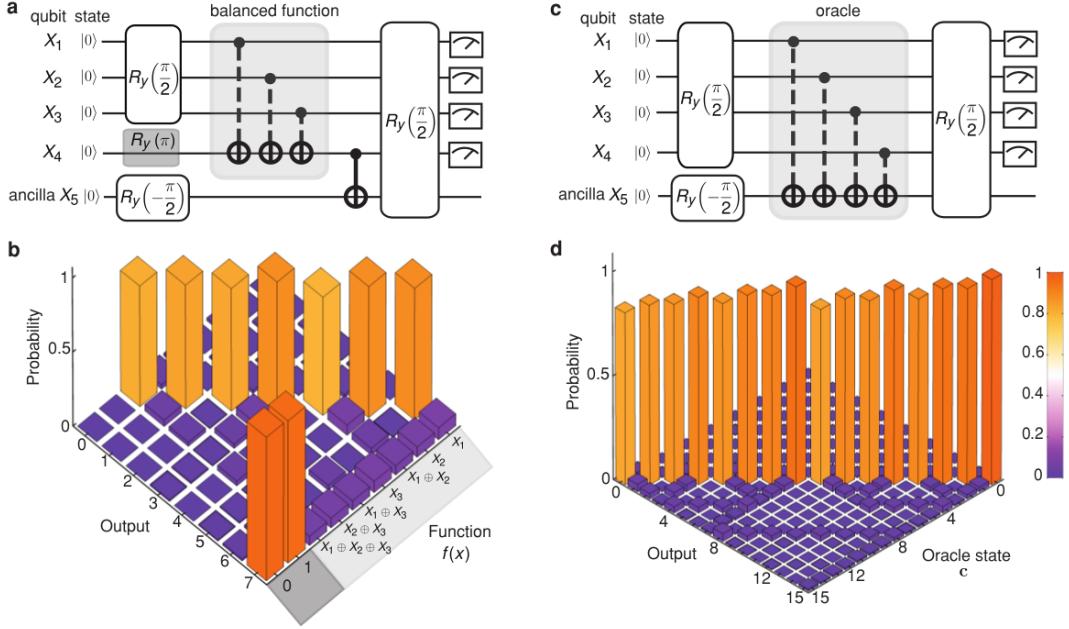


Figure 23: ([DLFL16]) A five-qubit implementation of the Deutsch-Josza (a) and Bernstein-Vazirani algorithm (b).

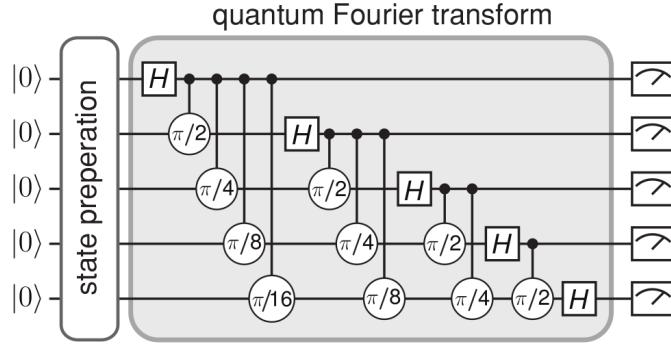


Figure 24: ([DLFL16]) A five-qubit implementation of the quantum Fourier transform.

nearest-neighbour Mølmer-Sørensen interactions [MS99] (XX-gates), a more sophisticated interaction for multi-particle entanglement which produces effective conditional dynamics in order to realize two-qubit gates such as CNOT. Moreover, the Raman beat-notes are tuned close to resonance ν_0 , yet slightly detuned down to $\nu_0 \pm \nu_x$ by an offset ν_x , in order to induce the necessary coupling. In addition to our previous discussion on single-qubit and two-qubit gates, Debnath et al. used a modern variant of MS-interactions by decomposing a CNOT gate into geometric phase gates, an approach that preserves the action of the CNOT, yet allows for an efficient and less error-prone realization by using the collective motion of the chain [MS99][HRB08].

In an attempt to explore novel architectures towards more scalable ion traps, Fellek et al. [FHM16] at the Georgia Institute of Technology have independently implemented the Bernstein-Vazirani algorithm with a chain of $^{171}\text{Yb}^+$ qubits using ion-transport in a microfabricated planar surface trap. Similarly, the algorithm succeeded at determining the unknown string with a success probability of 97.6%, for two ions, and 80.9% in the case of three ions, using only a single oracle query. The gate implementations and optics are similar to the Maryland setup: Single qubit gates are performed using a laser at 355nm wavelength driving Raman transitions. Two-qubit gates are also provided by nearest-neighbour Mølmer-Sørensen interactions [MS99].

10.2.7 Decoherence and Sources of Error

In this section, we discuss decoherence mechanisms and error sources that drive imperfections in trapped ion quantum computation. Just like in classical computation, a bit flip error $|g\rangle \leftrightarrow |e\rangle$ in a quantum state is devastating. As in most quantum devices, these errors typically occur during population inversion of the energy eigenstates of the system due to photon absorption or spontaneous emission and propagate through all subsequent computations. Thus, in any qubit manipulation requiring the use of lasers, there is some probability of driving an unwanted transition to other electronic levels.

In practice, alkali-earth-like metals, such as $^{171}\text{Yb}^+$, exhibit lifetimes of metastable states at about $2 - 3\text{s}$ [OYM07], hence the coherence time of hyperfine states is several orders of magnitude longer than the gate times ranging at microseconds. This fact makes ion-trap quantum computers surprisingly resistant to memory errors. On the contrary, external charge fluctuations in superconducting devices suffer considerably from bit-flip errors.

In terms of operational errors, both the original Cirac Zoller proposal and the XX-gates using nearest-neighbour Mølmer-Sørensen interactions are highly affected by population changes due to *motional heating*. Electromagnetic background radiation at the trap frequencies can create motional quanta that subsequently corrupt two-qubit operations. For example, due to the collective motional degree of freedom of the ion chain and the need for strict ground-state cooling, any two-qubit operation in the Cirac Zoller proposal is significantly prone to highly correlated errors as a result of spontaneous emission or electromagnetic radiation. This suggests that independent noise models, such as in Section 4.8, are non-physical in the context of an ion-trap architecture.

Phase flip errors are more subtle and have important fatal consequences in most quantum computations, as demonstrated in the well known *Ramsey experiment*. The phase evolution of a hyperfine state depends strongly on the magnetic field. A superposition of two state evolves due to individual magnetic moments and thus experiences dephasing due to energy fluctuations resulting from a fluctuating magnetic field. Typically, phase flips occur if the rf-Paul trap exhibits voltage fluctuations at the trap electrodes. In fact, the coherence time of ion-trap quantum computers is currently mostly limited by magnetic field fluctuations in the order of just a few milliseconds [HRB08]. Due to the fact that $^{171}\text{Yb}^+$ produces qubits with the same magnetic moment $m_F = 0$, the ytterbium ion is a popular choice to reduce dephasing effects.

In order to address errors during large scale computations in the future, one strives to adopt error correction, a procedure we discussed in Section 4.9. As one increases the number of qubits in a linear ion-trap architecture, i.e. the size of the ion chain, the addressing of individual ions with focused lasers onto the chain becomes increasingly difficult and complicates two-qubit operations with additional practical sources of error. Moreover, growing mass of the ion chain also results in reduced coupling on the sideband transitions through the Lamb-Dicke parameter.

In the next section, we shed light on the overall performance of the ion-trap architecture, as compared to a solid-state device.

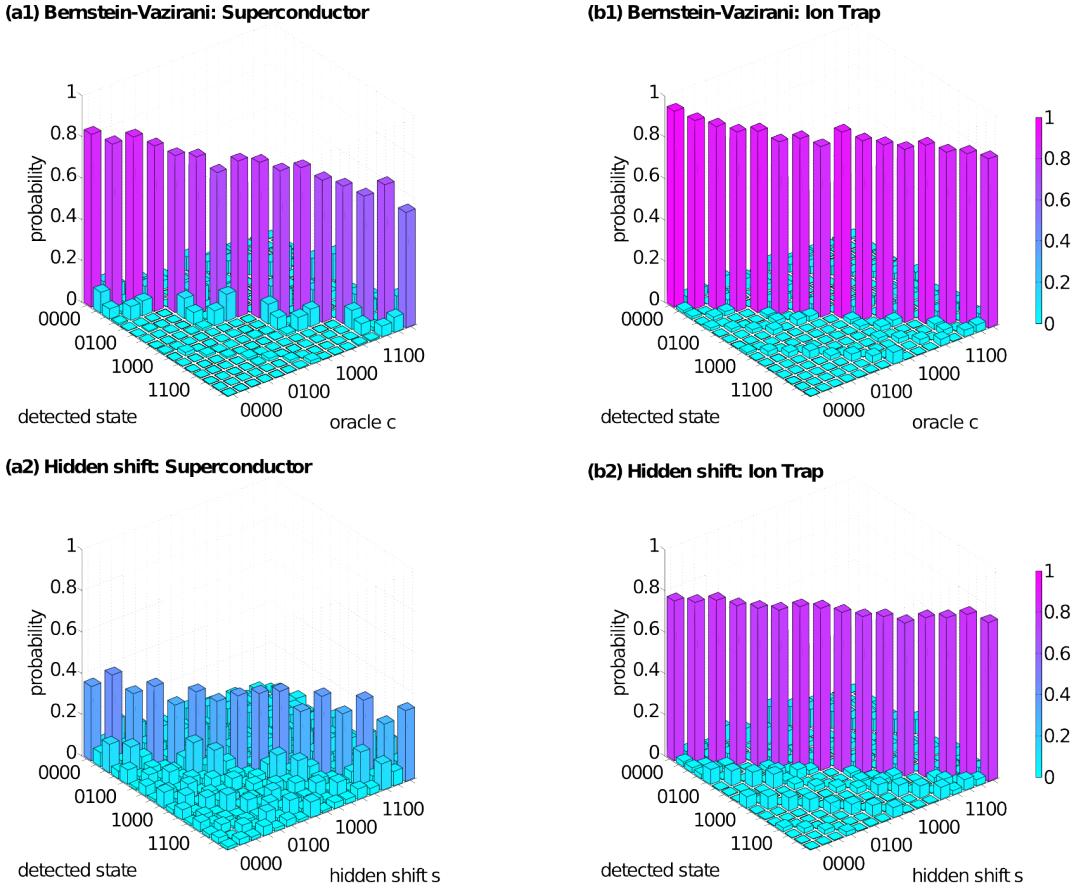


Figure 25: ([LMRD17]) Performance comparison between an ion-trap quantum computer compared to a solid-state architecture when running the Bernstein-Vazirani algorithm and Simon’s hidden-shift algorithm.

10.3 Performance Comparison of Quantum Computing Architectures

In the previous sections, we described how new advances in quantum computing technology made it possible to program algorithms from a high level user interface. To this date, particularly the ion-trap and solid-state architectures have reasonably grown in maturity, allowing for a variety of standard algorithms to be tested and evaluated for performance. Recently, IBM has launched a public-access demonstration of a five qubit superconducting device that can be run via their *Quantum Experience* cloud service.⁴ Building up on previous work at the University of Maryland [DLFL16], Linke et al. [LMRD17] from the Monroe group put forward a state-of-the-art comparison between the two leading quantum technologies, its own local ion-trap implementation vs. IBM's superconducting transmons, using the *Quantum Experience* platform. It was shown that, overall, the ion-trap quantum computer currently achieves higher success probabilities over the solid-state implementation from the IBM platform. Average success probabilities for running the Bernstein-Vazirani algorithm ranged around 85.1% for ion-traps, and around 72.8% for the superconducting device [LMRD17], see Figure 25. Particularly concerning gate times, noticeable differences were observed. Typical ion-trap gate times for single-qubit operations for an ion-trap computer averaged at around $20\mu\text{s}$ and $250\mu\text{s}$ for two-qubit gates, while superconducting circuits reached times of only 130ns and around 250 – 450ns, respectively. Overall, while current solid-state devices feature vastly higher clock-speeds, the ion-trap currently shows substantially higher absolute fidelities and longer coherence times. Nevertheless, the runner-up technology of the solid-state architecture offers a substantial promise for scalability. It remains to be seen which of the two architectures, if any, is going to establish itself as the leading scalable quantum computing technology of the future.

11 Conclusion

In this thesis, we shed light on how quantum algorithms achieve promising speed-ups over classical algorithms in the context of computational learning theory, even in the presence of noise. As quantum computational supremacy has yet to be demonstrated for a well-defined problem using only a few noisy qubits, the study of quantum learning remains a particularly relevant area of research. For further reading on the current status of quantum computational supremacy, we refer to an article by Harrow and Montanaro [HM17]. For an overview of recent progress in quantum learning theory, we refer to the survey by Arunachalam and de Wolf [AdW17].

By investigating the limitations of quantum algorithms through the use of blinding, we proposed suitable constructions for new notions of security under non-adaptive quantum chosen-ciphertext attacks. The pursuit of useful quantum-secure classical encryption schemes remains one of the key challenges in post-quantum cryptography. Therefore, further research is now needed to investigate whether classical communication is ultimately feasible in a quantum world. Finally, for further reading on the current status of post-quantum cryptography, we refer to a recent article by Bernstein and Lange [BL17].

⁴The IBM quantum experience platform can be found at <https://www.research.ibm.com/ibm-q/>

12 Open Problems

Due to the fact that **LWE** is easy once algorithms receive quantum superposition access to noisy samples on the secret string, it is tempting to explore circumstances in cryptography under which such access is granted. While it is reasonable to assume that providing such quantum access is ill-advised for public-key cryptography, it remains an open problem whether there are other realistic scenarios. One possible direction to investigate is *program obfuscation*, a recent breakthrough in cryptography which concerns the process of obscuring software or code in order to preserve functionality, yet hide sensitive information on the program itself with at most polynomial slowdown. An attacker in possession of a quantum computer could, in principle, implement the obfuscated circuit and then query it in superposition. Since *indistinguishability obfuscation* allows us to turn symmetric-key encryption schemes into public-key encryption schemes [SW14], this could potentially open up a door to providing quantum access to **LWE** samples in the context of the **LWE-SKES**.

Many of the classical notions of security are still widely unexplored in a quantum world, making both quantum cryptography and classical quantum-safe cryptography a highly relevant field of research. In 2016, Gagliardoni, Hülsing and Schaffner [GHS16] provided security standards of quantum indistinguishability under quantum chosen-plaintext-attacks and proposed secure quantum encryption schemes for which such security notions can be achieved. Further research is now needed to investigate secure constructions under a quantum chosen-ciphertext attack, in particular regarding a quantum indistinguishability notion of **QIND-QCCA1**. This indistinguishability setting naturally generalizes the **QCCA1** learning phase we considered in this thesis to a now fully quantum challenge phase in which the challenge ciphertext is also a quantum state. Thus, in the **QIND-QCCA1** indistinguishability game, both the learning phase and the challenge phase concern fully quantum communication. While many classical constructions for **CCA2** security already exist, it also remains an open problem to define a fully quantum notion of indistinguishability for **QCCA2** that allows for secure quantum encryption schemes.

A Supplementary Material

Lemma A.1 ([Wil13]). *Let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be quantum states and consider any POVM $\mathcal{E} = \{E_i\}_{i \in I}$. We define $p_i = \text{tr}[\rho E_i]$ and $q_i = \text{tr}[\sigma E_i]$ as the corresponding probability distributions over measurement outcomes labeled by $i \in I$. Then, the statistical distance between the distributions p_i and q_i is upper bounded by the trace distance between ρ and σ :*

$$\delta(p_i, q_i) \leq \delta(\rho, \sigma). \quad (\text{A.1})$$

Proof. Since $\rho - \sigma$ is Hermitian, we can use spectral decomposition to find a set of orthogonal vectors such that it is diagonalized:

$$\rho - \sigma = \sum_i \lambda_i |i\rangle \langle i|. \quad (\text{A.2})$$

Furthermore, consider splitting the vectors in terms of positive and negative eigenvalues by defining:

$$P := \sum_{\lambda_i \geq 0} \lambda_i |i\rangle \langle i|, \quad Q := \sum_{\lambda_i < 0} |\lambda_i| |i\rangle \langle i|. \quad (\text{A.3})$$

P and Q have orthogonal support and we can write $|\rho - \sigma| = |P - Q| = P + Q$, as the absolute value of a Hermitian operator takes the absolute value of its eigenvalues. Let $\mathcal{E} = \{E_i\}_{i \in I}$ be any POVM. Then, it follows that:

$$\begin{aligned} \delta(p_i, q_i) &= \frac{1}{2} \sum_i |\text{tr}[E_i \rho] - \text{tr}[E_i \sigma]| \\ &= \frac{1}{2} \sum_i |\text{tr}[E_i(\rho - \sigma)]| \\ &= \frac{1}{2} \sum_i |\text{tr}[E_i(P - Q)]| \\ &\leq \frac{1}{2} \sum_i \text{tr}[E_i(P + Q)] \\ &= \frac{1}{2} \sum_i \text{tr}[E_i|\rho - \sigma|] \\ &= \frac{1}{2} \sum_i \text{tr}|\rho - \sigma| \\ &= \delta(\rho, \sigma), \end{aligned}$$

as the POVM \mathcal{E} requires that $\sum_i E_i = \mathbb{1}$. □

Lemma A.2 ([NC10]). *Let $\rho = |\psi\rangle \langle \psi|$ and $\sigma = |\phi\rangle \langle \phi|$ be pure states. Then the trace distance between ρ and σ can be expressed in terms of the fidelity:*

$$\delta(\rho, \sigma) = \sqrt{1 - F(\rho, \sigma)^2} = \sqrt{1 - |\langle \psi | \phi \rangle|^2}. \quad (\text{A.4})$$

Lemma A.3 ([KL15], Markov's inequality). *Let X be a nonnegative random variable. Then, for any $a > 0$:*

$$\Pr[X \geq a] \leq \frac{\mathbb{E}[X]}{a} \quad (\text{A.5})$$

References

- [AB08] Aharonov, D., and M. Ben-Or, *Fault-tolerant quantum computation with constant error rate*, SIAM Journal on Computing 38(4), pp. 1207-1282, 2008.
- [AdW17] Arunachalam, S., de Wolf, R., *A survey of quantum learning theory*, CoRR, abs/1701.06806, 2017.
- [BKW94] Blum, A., Furst, M.L., Kearns, M.J., Lipton, R.J., *Cryptographic Primitives Based on Hard Learning Problems*, Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 278-291. Springer, Heidelberg, 1994.
- [BKW03] Blum, A., Kalai, A., and Wasserman, H., *Noise-tolerant learning, the parity problem, and the statistical query model*, Journal of the ACM, 50(4):506-519, 2003.
- [BL17] Bernstein, D., J., Lange, T., *Post-quantum cryptography*, Nature 549, 188-194, doi:10.1038/nature23461, 2017.
- [BV93] Bernstein, E., Vazirani, U., *Quantum Complexity Theory*, Proc. 25th ACM Symposium on Theory of Computation, pp. 11-20, 1993.
- [BZ13] Boneh, D., Zhandry, M., *Secure Signatures and Chosen Ciphertext Security in a Quantum Computing World*, Proceedings of CRYPTO, 2013.
- [CD10] Childs, A., van Dam, W., *Quantum algorithms for algebraic problems*, Reviews of Modern Physics 82, 1-52, 2010.
- [CSS14] Cross, A. W., Smith, G., Smolin, J. A., *Quantum learning robust against noise*, Physical Review A, 92(1):012327, 2015.
- [CZ95] Cirac, J. I., Zoller, P., *Quantum computations with cold trapped ions*, Phys. Rev. Lett. 74 (20), 4091-4094, 1995.
- [CZ97] Cirac, I., Zoller, P., Kimble, J., Mabuchi, H., *Quantum state transfer and entanglement distribution among distant nodes in a quantum network*, Phys. Rev. Lett. 78, 3221, 1997.
- [CZ01] Cirac, J. I., Duan, L. M., Zoller, P., *Quantum optical implementation of quantum information processing*, Lecture notes, "Experimental Quantum Computation and Information" Proceedings of the International School of Physics "Enrico Fermi" (2001), Course CXLVIII, p. 263, arXiv:quant-ph/0405030v1
- [Deu85] Deutsch, D., *Quantum theory, the Church-Turing principle and the universal quantum computer*, Proceedings of the Royal Society of London A400, pp. 97-117, 1985.
- [Deu89] Deutsch, D., *Quantum Computational Networks*, Proceedings of the Royal Society of London A425:73, 1989.
- [DHL05] Donley, E. A., Heavner, T. P., Levi, F., Tataw, M. O., Jefferts, S. R., *Double-pass acousto-optic modulator system*, Rev. Sci. Inst. 76, 063112, 2005.
- [DiV00] DiVincenzo, D., *The Physical Implementation of Quantum Computation*, Fortschritte der Physik, Volume 48, Issue 9-11 September, pages 771-783, 2000.
- [DJ92] Deutsch, D., Josza, R., *Rapid solution of problems by quantum computation*, Proc. R. Soc. Lond. A 1992 439 553-558; DOI: 10.1098/rspa.1992.0167, 1992.
- [DLFL16] Debnath, S., Linke, N. M., Figatt, C., Landsman, K. A., Wright, K., Monroe, C., *Demonstration of a small programmable quantum computer with atomic qubits*, Nature 536, 63-66, (04 August 2016) doi:10.1038/nature18648, 2016.

[Fey82] Feynman, R., *Simulating physics with computers*, Int. J. Theoret. Phys. 21, 467, 1982.

[FHM16] Fallek, S. D., Herold, C. D., McMahon, B. J., Maller, K. M., Brown, K. R., Amini, J. M., 1,5 *Transport implementation of the Bernstein-Vazirani algorithm with ion qubits*, New J. Phys. 18 (2016), doi:10.1088/1367-2630/18/8/083030, 2016.

[GGM86] Goldreich, O., Goldwasser, S., Micali, S., *How to construct random functions*, Journal of the ACM , 33(4):792-807, 1986.

[GHRL03] Gulde, S., Häffner, H., Riebe, M., Lancaster, G., Becher, C., Eschner, J., Schmidt-Kaler, F., Chuang, I. L., Blatt, R., *Quantum information processing with trapped Ca^+ ions*, Proc. R. Soc. Lond. A 361, 1363-1374, 2003.

[GHS16] Gagliardoni, T., Hülsing, A., Schaffner, C., *Semantic security and indistinguishability in the quantum world*, in Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III, pages 60-89, 2016.

[GK17] Grilo, A., Kerenidis, I., *Learning with Errors is easy with quantum samples*, as appeared in QCRYPT'17, Cambridge, UK, 2017.

[Gol04] Goldreich, O., *Foundations of Cryptography: Volume 2*, Cambridge University Press, New York, NY, USA, 2004.

[Gro96] Grover, L., *A fast quantum mechanical algorithm for database search*, In STOC, pages 212-219, ACM, 22-24, 1996.

[HH00] Hales, L., Hallgren, S., *An improved quantum Fourier transform algorithm and applications*, Proceedings of the 41st IEEE Symposium on Foundations of Computer Science, pp. 515- 525, 2000.

[HM17] Harrow, A., W., Montanaro, A., *Quantum computational supremacy*, Nature 549, 203-209, doi:10.1038/nature23458, 2017.

[HRB08] H. Häffner, C.F. Roos, R. Blatt, *Quantum computing with trapped ions*, Phys. Rep. 469, 155-203, 2008.

[Kit95] Kitaev, A. Y., *Quantum measurements and the Abelian stabilizer problem*, arXive e- print quant-ph/9511026, 1995.

[Kit97] Kitaev, A. Y., *Quantum operations: Algorithms and error correction*, RMS: Russian Mathematical Surveys, 52(6):1191-1249, 1997.

[KL15] Katz, J., Lindell, Y., *Introduction to Modern Cryptography*, CRC Press, 2nd Edition, 2015.

[KLLP16] Kaplan, M., Leurent, G., Leverrier, A., Naya-Plasencia, M., *Breaking Symmetric Cryptosystems using Quantum Period Finding*, CRYPTO, Springer, pages 207-233, doi=10.1007/978-3-662-53008-5-8, 2016.

[LOJC05] Langer, C., Ozeri, R., Jost, J. D., Chiaverini, J., Demarco, B., Ben-Kish, A., Blakestad, R. B., Britton, J., Hume, D. B., Itano, W. M., Leibfried, D., Reichle, R., Rosenband, T., Schaetz, T., Schmidt, P. O., Wineland, D. J., *Long-lived qubit memory using atomic ions*, Phys. Rev. Lett. 95, 060502, 2005.

[LMRD17] Linke, N. M., Maslov, D., Roetteler, M., Debnath, S., Figatt, C., Landsman, K. A., Wright, K., Monroe, C., *Experimental Comparison of Two Quantum Computing Architectures*, doi: 10.1073/pnas.1618020114, PNAS March 28, 2017, vol. 114 no. 13 3305-3310, 2017.

[MMKW99] Monroe, C., Meekhof, D. M., King, B. E., Itano, W. M., Wineland, D. J., *Demonstration of a fundamental quantum logic gate*, Phys. Rev. Lett. 75 (25), 4714-4717, 1995.

[Mon13] Monroe, C., *Quantum Simulations with Trapped Atomic Ions*, Slides from the Varenna Summer School on Ion Traps and Quantum Simulation, 2013.
http://iontrap.umd.edu/wp-content/uploads/2014/01/Lecture1_QSIM.pptx, (04.11.2017).

[MS99] Mølmer K., Sørensen A., *Multiparticle entanglement of hot trapped ions*, Phys. Rev. Lett. 82:1835-1838, 1999.

[NC10] Nielsen, M. A., Chuang, I. L., *Quantum Computation and Quantum Information*, Cambridge Univ. Press, Cambridge, 9th ed., 2010.

[OYM07] S. Olmschenk, K. C. Younge, D. L. Moehring, D. Matsukevich, P. Maunz, C. Monroe, *Manipulation and Detection of a Trapped Yb^+ Ion Hyperfine Qubit*, Phys. Rev. A 76, 052314, Published 19 November, 2007.

[Pau90] Paul, W., *Electromagnetic traps for charged and neutral particles*, Rev. Mod. Phys., vol. 62, pages 531-540, 1990.

[POF58] Paul, W., Osberghaus, O., and Fischer, E., *Ein Ionenkäfig*, Forschungsberichte des Wirtschafts- und Verkehrsministeriums Nordrhein-Westfalen 415, Westfälischer Verlag, 4, 15, 1958.

[Pre98] Preskill, J., *Reliable quantum computers*, Proceedings of the Royal Society A 454, pp. 385-410, eprint quant-ph/9705031, 1998.

[Reg05] Regev, O., *On lattices, learning with errors, random linear codes, and cryptography*, Journal of the ACM, 56(6):34, 2009. Preliminary version in STOC'05.

[Reg09] Regev, O., *The Learning with Errors problem*, Invited survey in CCC 2010, www.cims.nyu.edu/~regev/papers/lwesurvey.pdf, 2009.

[RS62] Rosser, J. Barkley; Schoenfeld, Lowell, *Approximate formulas for some functions of prime numbers*, Illinois J. Math. 6 (1): 64-94, 1962.

[Sha48] Shannon, C., *A Mathematical Theory of Communication*, Bell System Technical Journal. 27 (3): 379-423. doi:10.1002/j.1538-7305.1948.tb01338.x, 1948.

[Sho94] Shor, P. W., *Algorithms for quantum computation: discrete logarithms and factoring*, Proceedings of the 35th Annual Symposium on Foundations of Computer Science, Santa Fe, NM, Nov. 20-22, IEEE Computer Society Press, pp. 124-134, 1994.

[Sho95] Shor, P. W., *Scheme for reducing decoherence in quantum computer memory*, Phys. Rev. A 52, R2493(R), 1995.

[Sim97] Simon, D.R., *On the power of quantum computation*, SIAM journal on computing 26(5), 1474-1483, 1997.

[SW14] Sahai, A., Waters, B., *How to use indistinguishability obfuscation: Deniable encryption, and more.*, STOC, 2014.

[WBB03] D. J. Wineland, M. Barrett, J. Britton, J. Chiaverini, B. DeMarco, W. M. Itano, B. Jelenkovic, C. Langer, D. Leibfried, V. Meyer, T. Rosenband, T. Schätz, *Quantum information processing with trapped ions*, Phil. Trans. Royal Soc. London A, vol. 361, no. 1808, pages 1349-1361, 2003.

[WI79] Wineland, D. J. and Itano, W. M., *Laser cooling of atoms*, Phys. Rev. A, 20, 1521-1540, 1979.

[WMI98] D.J. Wineland, C. Monroe, W.M. Itano, D. Leibfried, B.E. King, D.M. Meekhof, *Experimental issues in coherent quantum-state manipulation of trapped atomic ions*, J. Res. Natl. Inst. Stand. Technol., vol. 103, pages 259-328, 1998.

[Wil13] Wilde, M., *Quantum Information Theory*, Cambridge University Press, 9781139525343, 2013.

[WZ82] W. K. Wootters, W. H. Zurek, *A single quantum cannot be cloned*, Nature 299, 802 - 803 (28 October 1982); doi:10.1038/299802a0, 1982.

[Zha12] Zhandry, M., *How to construct quantum random functions*, In Proceedings of FOCS, 2012. Full version available at the Cryptology ePrint Archives: <http://eprint.iacr.org/2012/182/>, 2012.

[Zha16] Zhandry, M., *A Note on Quantum-Secure PRPs*, arXiv:cs.CR/1611.05564v2, 2016.

Erklärung:

Ich versichere, dass ich diese Arbeit selbstständig verfasst habe und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, den 13.11.17