Ideal quantum protocols in the non-ideal physical world

Vedran Dunjko

Submitted for the degree of Doctor of Philosophy

Heriot-Watt University

School of Engineering and Physical Sciences

Institute of Photonics and Quantum Sciences

October 2012

Abstract

The development of quantum protocols from conception to experimental realizations is one of the main sources of the stimulating exchange between fundamental and experimental research characteristic to quantum information processing. In this thesis we contribute to the development of two recent quantum protocols, Universal Blind Quantum Computation (UBQC) and Quantum Digital Signatures (QDS). UBQC allows a client to delegate a quantum computation to a more powerful quantum server while keeping the input and computation private. We analyse the resilience of the privacy of UBQC under imperfections. Then, we introduce approximate blindness quantifying any compromise to privacy, and propose a protocol which enables arbitrary levels of security despite imperfections. Subsequently, we investigate the adaptability of UBQC to alternative implementations with practical advantages. QDS allow a party to send a message to other parties which cannot be forged, modified or repudiated. We analyse the security properties of a first proof-of-principle experiment of QDS, implemented in an optical system. We estimate the security failure probabilities of our system as a function of protocol parameters, under all but the most general types of attacks. Additionally, we develop new techniques for analysing transformations between symmetric sets of states, utilized not only in the security proofs of QDS but in other applications as well.

Acknowledgements

The research presented in this thesis has been performed under the guidance of my supervisors and mentors Elham Kashefi, Erika Andersson, and Gerald S. Buller. To them I extend my deepest gratitude for taking me on as their student, for their wisdom, knowledge and time which they shared with me. I could not have asked for better mentors.

I would like to thank the University of Edinburgh for providing me with a workspace at the School of Informatics, which was often very useful.

I have undoubtedly had a huge benefit in learning from my close collaborators, Tomoyuki Morimae, Anthony Leverrier and Robert. J. Collins. I cannot thank them enough for the selflessness, kindness and patience they extended to me during the many days and weeks we spent on endless discussions and paper writing. An extra thank you goes to Tomoyuki and Robert for letting me use and modify some of their illustrations in this thesis, Section 3 and Section 8, respectively. A big thank you is extended to Patrick Clarke and John Jeffers, with whom I worked on the quantum digital signatures project.

During the last year and a half, I had the opportunity to visit a number of other research groups and learn from wonderful people who are also superb scientists. Without exception, these visits resulted in elucidating discussions which profoundly influenced how I think about quantum information, physics and science in general. I would like to thank Damian Markham, one of the coolest physicist I know, for having me over in Paris and for many eye-opening discussions. Thank you Mário Ziman, Tomáš Rybár, Michal Sedlák and Daniel Nagaj for having me over in Bratislava, I had a great time! I am grateful to Terry Rudolph and Dan Browne for a wonderful visit to London, and to Richard Jozsa for inviting me to visit Cambridge and his group during this trip. Many thanks to David Gross who was kind enough to host me in Freiburg, I really enjoyed the discussions. A big thank you to Hans Briegel who welcomed me in Innsbruck where I also met Barbara Krauss, Wolfgang Dür, Julio De Vicente and many other wonderful people. I have enjoyed Innsbruck and all the discussions immensely. To Dominique Unruh, whom I met in Tartu I extend my sincere gratitude for his time, brainstorming sessions and for letting me mention some of his ideas in Sections 4 and 9 of this thesis. I had a superb time in Zürich! I extend my gratitude to Renato Renner for having me over, and a huge thank you goes to him, Matthias Christandl and their groups. That week was filled with vigorous discussions and for this I am very grateful. An extra thank you goes to Christopher Portmann and Fernando Brandão, for their involvement and discussions which ultimately led to the result presented in Section 4.3 of this thesis. I would also like to thank Shahram Mossayebi for valuable conversations, which influenced the writing of Sections 6 and 9. I am grateful to Einar Pius, Nikhil Ratanje, Daniel

Maldonado Mundo, Chaitanya Joshi and Adetunmise Dada, my “PhD siblings and cousins” for good times and useful discussions.

My approach to science was thoroughly influenced by my mentors during my education in Zagreb, Saša Singer, Sanja Singer, and Tomislav Domazet-Lošo, and for this I will always be grateful. You took me on as your student, and yet managed to be my friends as well. An additional thank you goes to Saša and Sanja for the discussions which eventually played a part in the resolutions of problems addressed in Section 10.

I do not know how to thank my family and my friends appropriately, as words seem insufficient. You are the people who made me – me, and all I can say is thank you for that.

Finally, I would like to thank Alida. I could not imagine the last few years without you, and I do not wish to. Thank you for letting me be a part of your life.

Contents

List of Publications by the Candidate

Parts of this thesis are based on materials which are published or in submission process. Below is the list of relevant publications, along with a brief commentary clarifying the contribution of the Candidate to each publication.

1. Vedran Dunjko and Erika Andersson
   Truly noiseless probabilistic amplification
   Physical Review A, Oct 2012 (accepted).

2. Vedran Dunjko and Erika Andersson
   Transformations between symmetric sets of states
   Journal of Physics A, 45(36):365305, Sept 2012 .

3. Tomoyuki Morimae, Vedran Dunjko and Elham Kashefi
   Ground state blind quantum computation on the AKLT state
   Jul 2012 (submitted).

4. P. J. Clarke, R. Collins, V. Dunjko, E. Andersson, J. Jeffers, and G. S. Buller
   Experimental demonstration of quantum digital signatures
   Jun 2012 (submitted).

5. Vedran Dunjko, Elham Kashefi, and Anthony Leverrier
   Blind quantum computing with weak coherent pulses
   Physical Review Letters, 108:200502, May 2012.

6. Vedran Dunjko and Elham Kashefi
   Algebraic characterisation of one-way patterns
   In DCM Proceedings Sixth Workshop on Developments in Computational Models: Causality, Computation, and Physics , pages 85-100, Jul 2010.

In publication 4, the Candidate has, along with EA and JJ, contributed to the approaches used in the security proofs. The Candidate suggested some of the experiments performed, and has provided the security proofs themselves. In publications 1,2,3,5,6 the Candidate produced the majority of the presented proofs and contributed to the developments of initial ideas. Publications 2 and 5 stemmed from the early posed, broad question whether Universal Blind Quantum Computation can work with continuous variables systems, originating from discussion between EA and EK. The initial idea in publication 3 was proposed by TM.

All the publications presented are joint work of all the authors, and a result of a true collaborative effort.

Introduction

UBQC and QDS – development of novel quantum protocols

Quantum information processing is an interdisciplinary field combining ideas from physics, computer science and information theory. The central promise of this field is that interesting and useful information processing tasks can be done by exploiting quantum mechanics. This field is an exciting intellectual playground, and many results often seem surprising as we still understand so little about the quantum world. Naturally, it is not the case that we are short on facts about this hidden quantum reality. However, it is a simple truth that *classical* behaviour – the behaviour of the world we are born into, and experience every day – is very different from *quantum* behaviour. So different that our innate intuition easily fails us in the latter case.

It is, therefore, not surprising that the development of quantum protocols, to actual physical and practical implementations, which exist in labs and eventually our daily lives, is full of unexpected impediments and side results.

The key questions of quantum information processing go well beyond the clear-cut approaches of finding new useful and practical things we can do by using quantum effects. The quantum reality is far from understood, and the investigation of quantum protocols, and their development, directly influences our understanding of the "quantumness" itself. Throughout this thesis we will present stages of development of two relatively novel quantum cryptographic protocols – universal blind quantum computation and quantum digital signatures. While our primary goal is understanding if and how the functionalities of these protocols can be realized in the physical world, we have strived to observe what the answers we get say about (quantum) physics itself.

Universal blind quantum computation (UBQC) [1] is the more recent of the two protocols. UBQC is a two-party protocol which addresses the problem of delegating quantum computation. In particular, one considers a powerful quantum server (whom we often refer to as *Bob*), taking a part not unlike the currently existing massive computational superclusters.

In UBQC, the server is capable of performing measurement-based quantum computation (MBQC) over a particular family of resource states called graph states. MBQC is universal model of quantum computation, the basics of which we discuss presently.

This server can be accessed by a client who wishes to run a quantum computation on the server, while maintaining privacy about her data. However, the client is assumed to have only limited quantum capabilities. Using UBQC, such a client can perform an arbitrary quantum computation

on a server, while maintaining perfect privacy. The privacy is guaranteed provided she has access to a simple quantum device which can generate specific single qubit states. In this work we consider the first analysis of the privacy properties of UBQC once any type of imperfections are introduced. We establish a concept of approximate blindness, which quantifies the disturbance imperfect devices incur on the security properties. Then, we introduce a preparation protocol using which arbitrary levels of privacy can be achieved, even if the client is restricted to relatively simple, and imperfect devices. Following this, we investigate whether UBQC can be adapted to work with a model of quantum computation other than graph state MBQC. The alternative model of interest may have practical advantages as the computation can be protected from noise by keeping the system of the server in an energy ground state. Finally, we outline the future directions the research into this protocol may take, and what we can learn about physics through deeper understanding of this protocol. Throughout Part 1 of this thesis we will interchangeably refer to the client and the server as *Alice* and *Bob*, respectively.

Part 2 of this thesis investigates quantum digital signatures (QDS) [2], a cryptographic protocol proposed about a decade ago. QDS is a multi-party protocol comprising a sender Alice and multiple recipients. We will consider the case of two recipients called Bob and Charlie. Roughly speaking, QDS enables Alice to send messages such that the recipients can verify the origin and integrity of the messages. Crucially, the messages are also transferable, meaning that if one of the recipients confirms the validity of a particular message, so will the other recipients. As a part of the research presented in this thesis, we analyse the performance of a proof-of-principle QDS experiment, realized in an optical system based on the proposal in [3]. In particular, we estimate the failure probabilities of the protocol, based on experimental data, taking into account various types of possible attacks. Following this, we take a step back and investigate the role quantum digital signatures may take in practice in the emerging quantum-digital world. In this discussion we place our focus on the practicality of the quantum digital signatures scheme in terms of ease of implementation, required rounds of classical and quantum communication as a function of the desired security levels. Additionally we propose "tweaks" to the realized scheme both at the experimental level and by proposing novel ways to achieve the desired security properties.

We begin by presenting the outline of the thesis.

Outline of this thesis

Part 1 : Universal Blind Quantum Computation

Chapter 1 – Introduction to universal blind quantum computing
In the first chapter we give a general introduction to universal blind quantum computing and state the security guarantees this protocol offers. We finalize the chapter by highlighting some of the issues that should be resolved if this protocol is to work in the real world. These issues we address in the following chapters.

Chapter 2 – World according to Alice: *UBQC under imperfections*
In Chapter 2 we study what security statements about universal blind quantum computation

can be guaranteed in any *realistic* realization of the protocol. We introduce a notion of *approximate blindness* we use to quantify the damage realistic imperfections may cause to the privacy of the client. Following this, we introduce a *remote blind quantum state preparation* (pre-)protocol which can be used in conjunction with universal blind quantum computation. By utilising this pre-protocol a client with access to very simple and realistic devices (no more involved than what is required for practical quantum key distribution) can achieve arbitrary levels of privacy with a manageable overhead in the communication costs.

Chapter 3 – World according to Bob: ***UBQC using alternative resources***

In Chapter 2 we proposed means to make blind quantum computing more practical, but from the perspective of the server. In this chapter, we address the complementary question and attempt to make the protocol more reasonable for the server. In particular, we address the feasibility of blind quantum computing which relies on novel resources useful for measurement based quantum computing as an alternative to the established graph-state-based approach. We present a protocol for blind quantum computing over the Affleck-Kennedy-Lieb-Tasaki state. This resource state, well-studied in condensed-matter physics, arguably allows for a more robust computation, since, unlike with the graph states, it can be protected from decoherence by an energy gap. In the process of developing this protocol we expose an interplay between communication, protection of the computation by an energy gap, and cryptographic privacy.

Chapter 4 – Discussion: future of UBQC

In this chapter we discuss other new developments in universal blind quantum computation. We suggest how UBQC can be used as a novel approach of investigating the foundations of quantum mechanics through a quantum complexity theory lens. We present a new rigorous proof of blindness of the original protocol, and discuss alternative models of this novel cryptographic primitive. Additionally, we discuss the specific properties of measurement based computation, which is the framework UBQC is formulated in, that make blind quantum computing possible. Finally, we briefly compare blind quantum computing with the new and famous classical protocol for secure delegated computation. This protocol known as "fully homomorphic encryption", took about 30 years to develop.

Chapter 5 – Side result: ***generalized phase map decomposition***

In this chapter, which can be read outside the context of blind quantum computing, we analyse the relationship between a computational process characterised by a one-way computation over graph state, and the actual linear map this computation implements. In the process of this analysis we re-capture certain known results on which types of quantum computation can be efficiently classically simulated.

Part 2 : Quantum Digital Signatures

Chapter 6 – Introduction to quantum digital signatures

In this chapter we briefly introduce the standard basic cryptographic communication notions, such as private channels, authentic channels, and digital signatures. Following this, we explain the idea behind quantum digital signatures, and present a known proposal for the implementation of this protocol in optical systems.

Chapter 7 – Experimental set-up

The proposal for the experimental demonstration of signature distribution presented in Chapter 6 has been implemented at Heriot–Watt University. In this section, we present the experimental set-up, and describe the experimental results which we use for the estimation of the security properties of our system.

Chapter 8 – Security analysis of the experiment

In this chapter we present the precise security statements, and calculate the security parameters of our set-up, based on experimental data. The system is proven secure for all but the most general types of attacks. For the most general types of attacks we provide a plausibility argument that they should not jeopardize the security of the system, however the rigorous proofs are left for further research.

Chapter 9 – Discussion: future of QDS

In chapter 9 we discuss initial new ideas on how the presented quantum digital signature scheme may be improved. These include proposals for the advancement of our implementation of the system, but also novel approaches to the QDS problem which potentially offer substantial practical advantages. We finalize this chapter by comparing QDS to related proposals in classical cryptography.

Chapter 10 – Side results: *some properties of symmetric sets of states and applications*

In the final chapter we present a couple of theoretical results which stemmed from the research in QDS and UBQC. In particular, we characterise the properties of transformations between symmetric sets of coherent states, which appear in both universal blind quantum computing and our realization of quantum digital signatures. We give two applications of the theory presented. First, we consider the properties of transforms converting phase encoded coherent states into relative-phase encoded qubit states. These transforms could, for instance, be used in UBQC. We calculate optimal success probabilities, and give a proposal of how the optimal transform may be approached in linear optics, asymptotically. Following this, we address the problem of truly perfect amplification of coherent light, viewed as a transform between symmetric sets of states. We calculate the success probabilities, and present schemes for phase-dependent amplification with unlimited gain which can also, albeit sub-optimally, be realized using linear optics.

Contributions

The main contributions presented in this thesis can be grouped in three categories – UBQC-related, QDS-related, and Side results.

UBQC-related contributions In this thesis we have presented the first framework for approximate blindness which allows for the quantitative analysis of the privacy properties of a UBQC protocol when imperfections are present. Additionally, we have given a Remote Blind qubit State preparation protocol that allows a client to run a UBQC protocol with arbitrary levels of privacy using realistic devices. These results have been published in [4], and presented in Chapter 2. Following this, we have studied the feasibility of UBQC using substantially different resources on the side of the server in an attempt to make the server's required computation more robust. This study showed an interplay between cryptographic properties of privacy, robustness of the server's computation, and communication assumptions. These results are presented in Chapter 3 and in [5]. In Chapter 4 we have presented sketches of novel results addressing the connections between UBQC, quantum complexity theory, and foundations of quantum mechanics. Additionally, we have given a new rigorous proof of blindness of the original Protocol 4.3.1, and provided additional new intuitive arguments why blindness should hold in this protocol. These arguments clarify why UBQC can be presented in the framework of measurement-based quantum computation more naturally than in the standard circuit model (presented in 4.3.2). The manuscripts containing the details of the outlined results presented here are in preparation.

QDS-related contributions The main contribution of this thesis related to QDS is the security analysis of the performed experiment. Techniques we have used could be adapted in the security analyses of future experiments. The security analysis has been presented in Chapter 8, and exposed in detail in [6]. Additionally, in Chapter 9, we have outlined novel ways to improve on the existent QDS schemes. The manuscripts with the details of the results outlined here are currently in submission process and in preparation.

Side results In the course of research into UBQC and QDS, we have produced a number of related side results which are interesting also outside the context of the respective protocols. Thus, in Chapter 5 we have explored the structure of the map realized by the measurement-based quantum computation of the type used in UBQC. This structure has been shown to bear relevance for the problems determining universal resources in MBQC, but also the characterisation of measurement-based computations which can be classically simulated. These results have been published in [7]. During the security analysis of the QDS experiment, we have extensively used the "symmetricity"of the quantum states which appear in this protocol. Motivated by this, we have developed a more general theory of properties of transformations of such symmetric states, and we have suggested the application of the results to UBQC and the problem of amplification of coherent light. These results are presented in Chapter 10 and in [8, 9].

Parts of this thesis are based on materials appearing in published papers and papers which are in the submission process at this time. For the list of publications of the author, along with a commentary on the author's contribution to each publication, please refer to the section List of Publications by the Candidate following the Contents of this thesis.

A word on prerequisites

In writing this thesis we have assumed that the reader is acquainted with the basic notions in quantum information and quantum computation, along with the accompanying level of mathematics. Since Part II of this thesis deals with an experiment realized in an optical system, there we assume familiarity with the basics of quantum optics as well.

We have put every effort to keep the notation and terminology consistent with most standard literature, such as the famous Nielsen and Chuang's textbook on quantum information and computation [10], Leonhardt's primer on quantum optics [11] and Barnett's recent "Quantum Information", which touches the basics of both quantum information and quantum optics [12]. Other notation and concepts, which are less than standard, such as the measurement-based computation model, have been defined and introduced in the text as required.

Aside from the notions in quantum optics and quantum information, parts of this thesis deal with ideas in theoretical computer science, such as classical and quantum complexity classes. The basic notions concerning computational complexity and computability in general are to some extent presented both in [10] and [12], and for a more mathematical and extensive introduction we recommend Sipser's introductory textbook on the topic [13]. However, as some of the complexity classes we mention in this work go beyond introductory textbooks, we give a list of definitions of all the classes used throughout the presented work at the end of this thesis in Section 10.3.2.

Part I

Universal Blind Quantum Computation

Chapter 1

Introduction to universal blind quantum computation

Universal blind quantum computation (UBQC) is a quantum protocol which enables a client, who does not have quantum computing capabilities, to delegate her computation to a quantum server. This protocol guarantees that the client's input, algorithm and output remain hidden from the server provided the client is capable of producing random separable single qubit states.

1.1 Delegating quantum computation and privacy

Building quantum computers seems hard. Well over 27 years have passed [14, 15] since the ideas of quantum computers were first proposed and formalized. In the decades that followed, focused efforts by experimental and theoretical physicists alike have been invested towards the goal of building this magical device. These continuing efforts have overcome considerable obstacles and setbacks, and yet the objective of true universal scalable quantum computers still seems distant. However, to say no forward momentum was gained is false. The optimistic experimentalist now believes that one day, in the foreseeable future, a quantum computer will be built [1]. However, even for an optimist, the dream of small and privately owned quantum computers, akin to the ubiquitous laptops and desktops, remains well out of reach. Realistically, large quantum servers may in the near future take a role similar to that occupied by massive superclusters today. The need for delegated computation is gaining more and more momentum in our cyber-society. Large servers themselves are becoming de-localized and one of the buzzwords computer science uses with increasing frequency is "cloud computing". In cloud computing data storage and processing tasks services are offered to a heterogeneous community of end users, supplied by a complex scattered network of various types of machines. Optimistically, in the near future quantum computing centres will become available as an important component in the large information processing cloud, solving problems perhaps beyond the reach of classical systems. However, since the computational cloud will be accessed "online" and "remotely", the issue of the security and the privacy of the process becomes paramount. In the modern world it will not suffice to have a large number of clients access large quantum clusters using simple devices. What is required is that the clients also enjoy full privacy guaranteed by an *efficient cryptographic scheme*.

[1] An example of such an optimistic experimentalist is prof. Anton Zeilinger, who expressed his firm belief that quantum computers will be built during a seminar held at Heriot-Watt University in 2011.

We now review and compare some of the suggested solutions to the problem of delegating quantum computation with privacy, with focus on the resources required and the security guarantees of each proposal.

In 2005 Childs [16] considered "secure assisted blind quantum computation". In Child's proposal the server is a universal quantum computer, capable of performing arbitrary measurement and unitary evolution of its registers. The client is assumed to possess a limited quantum machine. In particular, she is capable of storing qubits (has quantum memory), can "re-route" qubits (which is equivalent to a SWAP gate), and can perform single qubit Pauli X and Z gates. A bi-directional quantum communication channel is assumed to exist throughout the run-time of the computation. Additionally, it is required that the client has a random number generator.

Since the client has quantum memory, and the capability to perform Pauli X, Pauli Z and the SWAP gate, to achieve computational universality, she only needs to use the server to perform a Hadamard, $\frac{\pi}{8}$ and the two-qubit controlled not (CNOT) gate. The central idea behind Child's proposal is to ask the server to perform the desired gates on her quantum bit(s) sequentially, but only after she had encrypted her qubits using a quantum one-time pad [17]. The quantum one-time pad comprises applying a random Pauli operator on each qubit. From the server's point of view, such a quantum state underwent a maximally depolarizing channel, and contains no information about the initial state of the client's qubits. However, since the client knows which gates (which Pauli operators) have been applied in encryption, she can always decode the state. Thus, to perform a Hadamard gate on a single qubit, the client applies a randomly chosen Pauli operator on the qubit, sends it to the server who applies the Hadamard gate, and returns the qubit. For the case of the two-qubit CNOT gates, the client sends the two qubits, each one time padded, and the server applies the CNOT gate, and returns the qubits.

At the face of it, the server performed these two operations on encrypted data, so the correctness of the procedure is not obvious. Here, Childs uses the relatively simple commutation relations between the Hadamard and CNOT gates, and the Pauli operators (stemming from the fact that the so-called Clifford group of operators, which contain CNOT and Hadamard stabilize the Pauli group, which we will discuss later). To obtain the correct computational output, the client will simply have to decrypt her qubits using a different (but still Pauli) operator than she used for encryption.

How these two gates are implemented by the help of the server we illustrate in Figure 1.1.

This enables the client to use the server to perform any Clifford gate on her register, without revealing any of her data. However, to achieve computation universality the client has to be capable of applying the $\frac{\pi}{8}$ gate on her register as well.

Since the $\frac{\pi}{8}$ gate is not in the so-called Clifford group, and the construction for the application of this gate with the help of the server is more involved two-round subroutine. In particular, the proper decoding of a quantum one-time padded qubit, upon which a $\frac{\pi}{8}$ was performed involves the $S = |0\rangle\langle 0| + i|1\rangle\langle 1|$ gate. However, this gate *is* a Clifford gate, so the client can have the server perform it for her, over encrypted data, while leaking no information. For details we refer

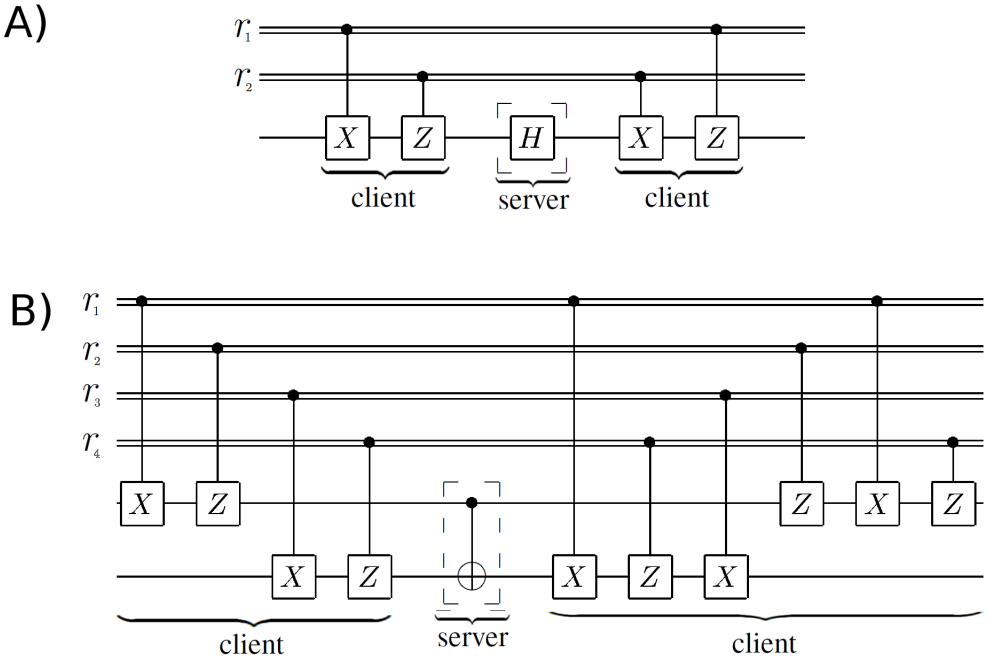

Figure 1.1. Server assisted application of the Hadamard and controlled-not gate. In image A), first the client applies a random Pauli operation on her qubit, which depends on the randomly chosen bits r_1 and r_2. The server applies the Hadamard gate on the encoded state, which is returned to the client and then decoded. Note that the encoding Z gate is decoded with the X gate and vice-versa due to the following commutation relations: $XH = HZ$ and $ZH = HX$. Image B) illustrates how the CNOT is implemented. Each qubit of the client's two-qubit state is quantum one-time padded by choosing random bits $r_1 \dots r_4$. The decoding circuit is again a result of the commutation relations of the controlled not gates and the Pauli gates (Figure adapted from [16]).

the reader to [16] but the end result is the same – the client gets the server to perform this gate on a qubit of her choice by sending one-time padded information to the server only. Similarly, Child's shows how the client can use the server to perform measurements.

While this approach intuitively seems to allow computation where the client does not leak any information about her data the exposition in [16] lacks a formal definition of security, and the proof it holds for the presented protocol. Ignoring this technicality, the security to the client is unconditional. However, the requirements on the client are relatively steep: an (arbitrary sized) quantum register, ability to perform SWAP and Pauli gates, and a two-way quantum channel accessible in the run time of the computation, and a random number generator. We note that the requirement for an arbitrary-sized quantum register may be relaxed a bit. Provided the client has the capacity to quantum-one time pad her private input and forward it to the server, she can perform all of her computation by manipulating only up to three qubits at a time. Thus, in principle, a constant-sized register could suffice.

In 2006 Arrighi and Salvail [18] proposed a different solution to the problem of delegated private quantum computation, by introducing the notion of *trap computations*. In their setting, the quantum server is capable of performing a function f on an input x. The function f is such that it can be realized as a unitary evolution of a register containing the input x, encoded in the state of the register $|x\rangle$, in the following way:

$$U_f|x\rangle|0\rangle = |x\rangle|f(x)\rangle.$$

The states $\{|x\rangle\}_x$ comprise an orthogonal basis we call the computational basis, with respect to all possible valid inputs x.

The client is assumed to be capable of efficiently generating random input-output pairs of the function of the form $(q, f(q))$, the quantum states $|q\rangle$ encoding the input q, but also the superposition state of any two valid inputs (q, q'): $1/\sqrt{2}(|q\rangle + i|q'\rangle)$. Such states are then called "quantum decoys".

The key idea behind the quantum decoys is to detect whether the server attempted to access information in the states the client sends: to learn the client's input $|x\rangle$, the server needs to measure this system in the computational basis. However, the decoy state is in a superposition of computational basis states, so such measurement can in principle be detected, as it will collapse the state. The basic protocol is presented in Protocol 1.

Protocol 1 Protocol of Arrighi and Salvail

1. The client generates selects a desired input x and a efficiently computes a pool of $2N$ randomly chosen valid input-output pairs (relative to the fixed function f) $\{(q, f(q))\}$
2. The client then prepares $N + 1$ states comprising N quantum decoy states of the form $1/\sqrt{2}(|q\rangle + i|q'\rangle)$ (where q and q' are inputs generated in the last step) and the state $|x\rangle$.
3. The client selects a state $|\psi\rangle$, uniformly at random, from the set of $N + 1$ states generated in the last step. She forwards this step to the server and asks him to perform U_f and return the output.
4. The server returns a state, which, if the server is honest, is $U_f|\psi\rangle|0\rangle$.
5. If the state $|\psi\rangle$ was a decoy state, the client checks for tampering by performing a measurement defined by the measurement operators

$$P_{intact} = \frac{1}{2}(|q; f(q)\rangle + i|q'; f(q')\rangle)(\langle q; fq| - i\langle q'; f(q')|).$$
$$P_{tamper} = \mathbb{1} - P_{intact}.$$

   If the result of this measurement corresponds to the measurement operator P_{tamper}, the client aborts. If the state $|\psi\rangle$ was the true input state $|x\rangle$, the client reads out the outcome.
6. If the pool is empty, the procedure halts, otherwise the client goes back to step 1.

The POVM element P_{intact} is just a projector onto the state an honest server would return. It depends on the pair q, q', but also the output values of the function f – $f(q)$ and $f(q')$. Because of this reason, the client is required to be capable of generating valid input output pairs for the desired function f. For this protocol Arrighi and Salvail prove the following security

theorem.

Theorem 1. *[18] Suppose the server has no prior information about the true input x. Let I be the mutual information of the server about the true input x at the end of the protocol. Let D be the probability the client detects tampering. Then, provided the server performs individual attacks only, we have:*

$$D \geq 1 - F(2^{I - log(|x|)})^N, \tag{1.1}$$

where F denotes the induced fidelity, and $|x|$ the size of the instance x.

Thus, for a fixed information gain by the server, the detection probability D approaches unity exponentially quickly with the number of decoys used.

Concerning the quantum requirements of the client, the protocol of Arrighi and Salvail dictates that the client can generate superpositions of computational basis states, and has a bi-directional quantum channel to the server. Additionally, the client has to be capable of performing the measurement described by the operators P_{intact}, P_{tamper}. This is a parametrized multi-qubit measurement which depends on the generated input-output pairs of the function f, as shown in Step 5 of Protocol 1.

Note that the client has to generate random input output pairs for the desired function f if Protocol 1 is to work. This restricts the classes of functions f which can be used to the class of so-called *random verifiable functions* which we present as defined in [18]:

Definition 1. *Let S and S' be two finite sets. A function $f : S \to S'$ is* random verifiable *iff for all N there exists an efficient probabilistic process which generates N input-output pairs $(q, f(q))$ such that the inputs (q's) are uniformly distributed in S.*

Some random verifiable functions are particularly interesting as they can be computed efficiently on the quantum computer, but are not known to be tractable for classical computers. However, it is not believed that all functions efficiently computable on quantum computers are random verifiable. Hence, the protocol of Arrighi and Salvai does not seem to allow the client to compute arbitrary BQP functions.

In 2010 (available on arXiv since 2008), Aharonov, Ben-Or, and Eban, were studying the concept of "quantum prover interactive proof systems" [19]. As we will elaborate later in Section 4.1, this complexity-theoretical concept also explores the question whether large-scale quantum mechanics can be falsified in practice, given that the simulation of large quantum systems seems intractable. In the derivations of their results they develop a delegated quantum computation scheme, in which a limited client can compute arbitrary quantum computation with the help of the server. A key requirement in their approach is that the client obtains a guarantee that the output of the computation is *correct*, a property which is called *verifiability*. While verifiability of delegated quantum computation is not the central research interest of this thesis, we will return to this property in Chapter 4. In [19] propose two protocols, in which a client, equipped with a constant size quantum computer, can compute arbitrary quantum computation with the help of an unbounded quantum server. In both protocols, the key idea is that the client uses Quantum Au-

thentication Schemes, which can be used for the authenticated exchange of quantum messages [17]. In the simple protocol, the client capable of performing universal quantum computation over some constant c qubits, encodes all the qubits she will require for her computation sequentially in an authentication code and sends them to the server. This encoding contains a quantum one-time pad, so all the data is inaccessible to the server. To perform individual (two-qubit) operations, the client requests the server to send back the desired qubits. The client decodes the qubits, and aborts if the quantum authentication decryption procedure detects an error. If no error is detected, the desired two qubit gate is applied by the client, the output re-encoded in a quantum authentication code, and sent back to the server.

In this simple protocol, the client simply uses the server as a large quantum information storage facility, and the use of quantum authentication codes ensures both data integrity and privacy. However, a bi-directional communication channel between the server and the client is required in the run-time of the protocol. The stronger protocol suggested in [19] is relatively involved compared to the proposals we have shown so far, so we will just present the basic idea.

To construct the second protocol, Aharonov, Ben-Or, and Eban show how by using the Polynomial Authentication Code [20], which is a type of a quantum authentication code, the server can perform any arbitrary quantum computation *over data encoded by the client*. To ensure the correctness of the computation the server needs classical communication with the client, and an ample supply of so-called Toffoli states. In particular, for Clifford computation, classical communication with the client is sufficient. To achieve computational universality (for which the addition of the Toffoli states will suffice), the server will additionally use Toffoli states, initially prepared by the client. More precisely in the second protocol, the client prepares the required quantum states for the computation. These include the desired data qubits, and the Toffoli states, and all is sequentially encoded in a Polynomial Authentication Code and sent to the server. For this the client requires a constant sized quantum computer. From this point on, only classical communication is needed. The client's privacy is unconditional.

In 2009 Broadbent, Fitzsimons and Kashefi proposed Universal Blind Quantum Computation (UBQC) [1], the protocol we will be shortly describing in detail. The framework they have used is that of Measurement-Based Quantum Computation (MBQC), in contrast to the more standard quantum circuit model used in the approaches of Childs and Aharonov, Ben-Or, and Eban.

In UBQC, the client is only required to produce single qubit states, randomly chosen from the set $\{1/\sqrt{2}(|0\rangle + exp(ik\pi/4)|1\rangle)\}_{k=0}^{7}$, and emit them to a more powerful quantum server. This preparation and the sending of qubit states can be performed off-line.

From these single qubit states, the server prepares an encrypted resource state which can be used for universal MBQC. During the actual run time of the computation, only two-way *classical* communication is required, along with simple arithmetical operations on the side of the client. By utilising the UBQC protocol the client can evaluate any (polynomial) quantum computation using the quantum server, and the privacy is unconditional.

Table 1.1 below summarizes the basic comparison of the performance of the approaches to delegating quantum computation we have mentioned:

Protocol	Childs	Arrighi and Salvai	Aharonov, Ben-Or and Eban	UBQC
Power of client	q. memory, Pauli gates, SWAP gate	preparing superpositions, multiqubit measurements	constant size q. computer	random single qubit states preparation
Functions which can be evaluated	BQP	random verifiable	BQP	BQP
Run-time quantum channel required	Yes	Yes	No	No
Security guarantee	information-theoretical (unproven)	efficient tampering detection (for indiv. attacks)	information-theoretical	information-theoretical

Figure 1.2. Comparison of delegated quantum computation protocols.

Delegating classical computation The idea of delegating difficult problems to large service centres, which are faster, better and larger than our small private devices, while maintaining the client's privacy has along history in classical computer science and cryptography.

Perhaps the first question relevant to secure delegated information processing was formalized as the problem of "computing over encrypted data". Already in 1978 Rivest, Adleman and Dertouzos [21] consider the problem of finding an encryption function with particular properties. An encryption function takes input form the set of plaintexts, the information we wish to process, to the set of ciphertexts, which are encrypted versions of our information. The desired property, suggested in [21], is the following: if the sets of plaintexts and ciphertexts are equipped with an algebraic structure — addition and multiplication, for instance — then the encryption function is required to be an algebraic structure homomorphism [2]. If a secure enough encryption function with this property existed, then anyone (in particular, a more powerful server) could perform large, time-consuming calculations over encrypted data, without decrypting it. From this, the client could reap the benefits of the calculation by direct decryption of the outcome, while maintaining privacy. The prototypical example of such a function presented in [21], is the RSA (Rives, Shamir, Adleman) public-key encryption function. This was the first example of a *partially* homomorphic encryption, as it is a multiplicative homomorphism only, and not a homomorphism of a computationally universal set of operations. Despite this promising initial result in 1978, and considerable interest in the problem, a fully homomorphic scheme was out of reach for the next 30 years. Finally, in 2009 Craig Gentry presented the first public-key based, computationally secure fully homomorphic encryption (FHE) scheme in one of the major conferences in the field. This result finally proved that, at least in principle, fully homomorphic encryption is possible [22]. While this result of Gentry and UBQC do not address the same problems, as both may have important roles to play in the future of delegated cloud computation, we will briefly address their mutual relationship later in Chapter 4.

[2]Roughly, a function from one structure to another is a structural homomorphism, if the function preserves the structures. A simple example are homomorphisms of, say multiplicative groups (which are simple algebraic structures). There we require that the homomorphic image of a product of two elements (in the domain) is the product of the homomorphic images of the two elements (in the range): $f(xy) = f(x)f(y)$.

Delegating computation and complexity theory In 1987 Abadi, Feigenbaum and Killian [23] studied the problem of computing over encrypted data from a complexity theoretical perspective. In their setting, they considered two players, A and B. Player A wants to compute the value $f(x)$, for some input x, but lacks the computational power to compute f. Player B hs the power to compute f for any input, is in principle computationally unbounded, and is willing to help player A. However, player A wishes to keep her input secret from player B. Thus any solution has to involve player A encrypting her data.

More formally, they consider an *encryption scheme* for function f which comprises and encryption function

$$E : Dom(f) \times \mathcal{K} \rightarrow Dom(f) \tag{1.2}$$

and a decryption function

$$D : Dom(f) \times \mathcal{K} \times Range(f) \rightarrow Range(f), \tag{1.3}$$

where $\mathcal{K}$ is a set of *keys*.

To compute $f(x)$ for a chosen instance x with the help of player B, player A computes $y = E(x, k)$ where k is chosen uniformly at random from $\mathcal{K}$. The value y is sent to player B, who then returns $f(y)$. Player A then computes $f(x) = D(x, k, f(y))$.

The message $y = E(x, k)$ could be correlated to the hidden input x and leak information about x to player B. To formalize this, Abadi et. use the language of random variables. Let X be the random variable taking values in the set of all possible inputs $\{x\}_x$ (the possible choices of player A), and let Y be the random variable taking values in the set of all encryptions for all keys k $\{y = E(x, k)\}_{x,k}$. Then we have the following definition:

Definition 2. *An encryption function E for f leaks (some function) L (of the input) if, for all a-priori distributions on X, for all statements z, the random variables X and Y are independent given $L(X) = z$.*

Intuitively, if E encrypts f while leaking L, then if $L(X) = z$ is revealed to player B before any message is sent from A, he learns nothing new once he receives y in the encryption scheme.

The idea above captures the basic notion of "computation with encrypted data". However, it can be generalized, by allowing intermediate communication and relaxing the necessity for a perfectly correct final outcome for player A. In a *generalized encryption scheme*, the players A and B are allowed a certain (polynomial in instance size) number of communication rounds m, and at each round, player A is allowed to perform a computation which is tractable for her. In the end, the absolute correctness of the final computed outcome for player A is relaxed as well: after m messages, A correctly guesses $f(x)$ with a bounded-error probability, say $2/3$. For more details on the reasoning behind bounded-error probability, please refer to the definition of the BPP class at the end of this thesis. The notion of *leaking something* is easily extended from the simpler case of an encryption scheme by considering not one random variable Y but m random variables corresponding to all the messages player A sends to B.

The questions of existence of generalized encryption schemes for interesting functions f in [23] centred on the relationship between the following three concepts:

- the computational power of player A, that is, how computationally hard are the functions E and D are allowed to be,
- the computational hardness of the function f player A wants to evaluate, and
- the amount and quality of information about x is allowed to be leaked to player B.

The computational hardness above is meant in complexity-theoretical sense.

Some relationships are obvious. For instance, if the encryption and decryption functions are allowed to be as hard as f, then nothing needs to be leaked to played B as player A can compute f on her own. Alternatively, if all information is allowed to be leaked to player B, player A needs no computational powers, as the encryption and the decryption function which satisfy this setting are the identity.

However, in [23] other relevant results are obtained. In particular, the following two theorems are proven.

Theorem 2. *[23] If there exists a generalized encryption scheme for a function f where the encryption and the decryption functions are in* ZPP*, such that no information about the input instance x is leaked to player B, then $f \in$ ZPP.*

The inverse of this claim is obvious: if $f \in$ ZPP and player A can evaluate ZPP functions, then player A can obtain $f(x)$ on her own, without revealing any information to player B.

A different result is obtained if the requirement on perfect secrecy is relaxed, as given by the following theorem.

Theorem 3. *[23] If there exists a generalized encryption scheme for an* NP*-hard function f, where the encryption and the decryption functions are in* ZPP*, such that only the* the size *of the instance x is leaked to player B, then the polynomial hierarchy (*PH*) collapses to the third level.*

In terms of the definition of leaking L given before, leaking the size means $L(x)$ is the size of x.

Complexity theory and cryptographic security have a long standing relationship. Almost the entirety of public-key-based cryptography is secure precisely under computational assumptions. Computation assumptions imply that certain functions are considered intractable (or overly time consuming) for any adversary, which is a complexity-theoretical statement. In contrast in [23], the security considered is information-theoretical (as player B is computationally unbounded), and the result above establishes a different flavour of links between complexity theory and cryptography.

Whether this result has potential bearing on UBQC we will briefly address in Chapter 4.

For the remainder of this introductory chapter, we focus on the fundamentals of UBQC.

The UBQC protocol is formulated in the framework of Measurement-based Quantum Computation (MBQC) [24, 25, 26, 27] and a thorough understanding of MBQC facilitates a simple way

to present the key ideas behind UBQC . For this reason, we will initially dedicate a bit of space to a particular model of MBQC relevant to UBQC.

1.2 Universal blind quantum computation

1.2.1 One-way quantum computation

In modern quantum information processing, the default model for studying quantum computation is the quantum circuit model (QCM). In the standard case of qubit-based QCM, an n-qubit computation is defined as the following process:

1. A *computational register*, an $n-$qubit system, is initialized to a pre-defined pure separable state, say $|0\rangle^{\otimes n}$.
2. A quantum programme is run on the register by *sequential application* of single and two-qubit gates (unitary evolutions of single or two-qubit subsystems).
3. The *classical* output of the computation is obtained by a prescribed measurement of one or more of the qubits of in the register.

Provided the set of gates the quantum machine can perform includes all single qubit gates, along with one (entangling) two-qubit gate, the process above is sufficient for realizing any arbitrary unitary evolution of the register. This property is called (quantum) computational universality. Most often, the two-qubit gate of choice is the controlled-NOT (CNOT) gate, but other gates such as controlled-Z gate ($\wedge Z$) achieve the same purpose. Moreover, the requirement for any arbitrary single-qubit gate can be relaxed by allowing only a small discrete set of (universal) gates, most often the Hadamard and the $\frac{\pi}{8}$ gate. Using this restricted set of single-qubit unitary gates, any arbitrary single qubit unitary can be *efficiently* approximated, which is a consequence of the Solovay-Kitaev Theorem [10]. With this restriction to a finite set of gates, the described quantum computation achieves *approximate* universality, which is sufficient for most intents and purposes, including the purposes of this thesis.

A crucial component in the circuit model is the on-demand two-qubit gate. Without it, any quantum circuit computation, defined as above, can not only be efficiently simulated on a classical computer, but is not even universal for classical computations [3]. The required two-qubit interactions such as the CNOT gate, central for computational power of the circuit model, have the capacity to generate entangled from separable states. The ability to generate entanglement suffices to achieve computational advantages over classical computers, but equivalent results may be obtained if pre-existing entanglement is used as a resource. Measurement-based quantum computation (MBQC) is one of many alternative models which achieve quantum computational universality. In MBQC the pre-existing entanglement in a large quantum resource state, rather

[3]The simulation of this process would comprise simple 2×2 matrix multiplications which are computationally easy. Additionally, without a method which allows one qubit to influence the state of another, *i.e.* a two-qubit gate, two bit logical gates, necessary for universal classical computation, such as XOR or NAND cannot be implemented.

than the capacity to generate entanglement in run-time, is a key element used to achieve quantum computational capabilities. Pre-existing entanglement has been known to be a valuable resource in general quantum information processing (QIP) tasks. From the beginnings of QIP, pre-existing entanglement (shared Bell-pairs, most often) has been used to achieve superdense coding [28], secure communication [29], and quantum teleportation [30]. In fact, the last functionality is intimately related to MBQC [31, 27]. In MBQC, the computation itself is driven not by local unitary gates, but rather by adaptive local measurements, as we presently explain.

For the purposes of understanding UBQC, we shall now quickly revise the relevant particular model of MBQC called the *one-way model*. In the one-way model the underlying resource is a *graph state* [32, 33], defined as follows:

Definition 3. *A* graph state $|G\rangle$ *corresponding to a graph* $G = (E, V)$*, where* E *and* V *are the sets of edges and vertices of* G*, respectively, is a pure quantum state given as:*

$$|G\rangle = \Pi_{(i,j)\in E} \wedge Z_{(i,j)} |+\rangle^{\otimes V}. \tag{1.4}$$

The single qubit quantum state $|+\rangle = 1/\sqrt{2}(|0\rangle + |1\rangle)$ is a uniform superposition of computational basis states $\{|0\rangle, |1\rangle\}$, and the operator $\wedge Z_{(i,j)}$ applies the two-qubit controlled-Z gate between the qubits i and j. The controlled-Z gate $\wedge Z$ is defined as:

$$\wedge Z = \sum_{b_0,b_1=0,0}^{1,1} (-1)^{b_0 \& b_1} |b_0 b_1\rangle\langle b_0 b_1|,$$

where with "&" we denote the bit-wise "and" operation. This operation applies the Pauli-Z operator on the second qubit if the first qubit was in the $|1\rangle$ state and does nothing if it was in the $|0\rangle$ state, and this defines it uniquely by linearity. The qubits in this resource state can be interpreted as vertices, and the entangling operations as edges, hence the name "graph states".

The computation itself is driven by single-qubit measurements as we now explain. In the model we shall consider, the measurements are determined by a single angle ϕ, and the corresponding observable (hermitian operator), denoted M^ϕ, is characterised by the (orthogonal) eigenvectors $|+_\phi\rangle := \frac{1}{\sqrt{2}}\left(|0\rangle + e^{i\phi}|1\rangle\right)$ and $|-_\phi\rangle := \frac{1}{\sqrt{2}}\left(|0\rangle - e^{i\phi}|1\rangle\right)$. The measurement outcome s (for "signal"), by convention attains values 0 and 1, corresponding to the possible post-measurement states $|+_\phi\rangle$ and $|-_\phi\rangle$, respectively.

For illustration purposes, consider the simplest non-trivial computation expressible in this model: a single qubit measurement over a 2 qubit chain, as illustrated in Figure 1.3.

The state of the system at cut A is $1/\sqrt{2}(|0\rangle|+\rangle + |1\rangle|-\rangle)$, which is a maximally entangled state. For the case of the 0 measurement outcome of the applied measurement M^ϕ, the state of the second qubit can be computed by applying the operator $\langle+_\phi| \otimes \mathbb{1}$ on the state in cut A [4]. Note that $\langle+_\phi| = \langle+|Z_{-\phi}$, where $Z_{-\phi} := |0\rangle\langle 0| + e^{-i\phi}|1\rangle\langle 1| = e^{i\phi Z/2}$. Thus we have the following

[4] This is equivalent to the application of the projector $|+_\phi\rangle\langle+_\phi| \otimes \mathbb{1}$, followed by tracing out of the measured qubit.

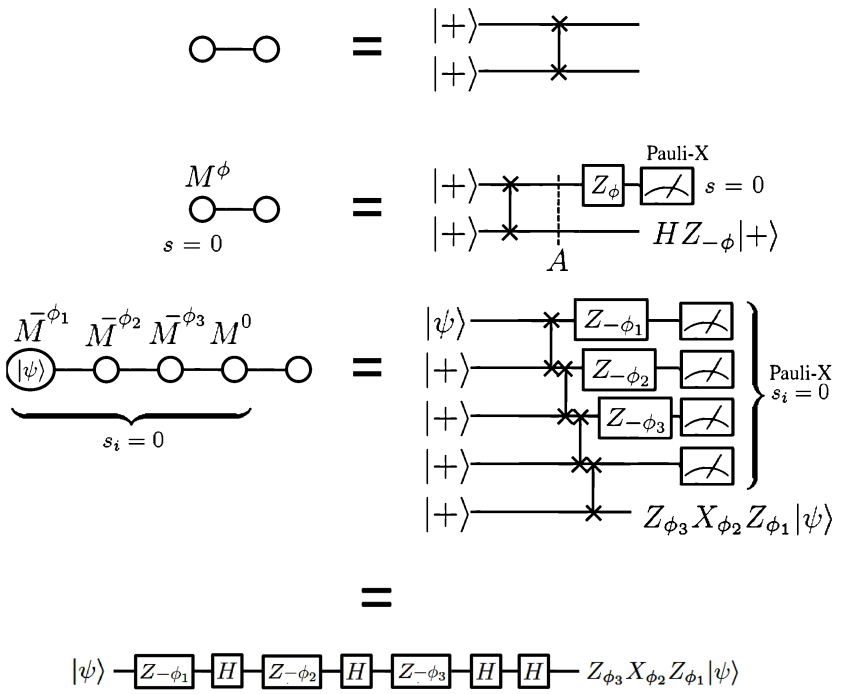

Figure 1.3. Projections and single qubit universality. On the left side of this figure, we present the standard notation used in MBQC. The right hand side illustrated the "extended circuit" – a direct translation of MBQC terminology to the circuit model. The chain of 5 qubits can be used to perform an arbitrary single unitary operation on an input single qubit state as illustrated in the third circuit from above. The final equality shows how measurement based quantum computation over a one dimensional graph state can be interpreted as a sequence of single qubit gates over a qubit line in the circuit model.

derivation, ignoring normalization factors:

$$(\langle +|Z_{-\phi} \otimes \mathbb{1})\left(|0\rangle|+\rangle + |1\rangle|-\rangle\right) = (\langle +| \otimes \mathbb{1})\left(|0\rangle|+\rangle + e^{-i\phi}|1\rangle|-\rangle\right) = \tag{1.5}$$

$$|+\rangle + e^{-i\phi}|-\rangle = HZ_{-\phi}|+\rangle. \tag{1.6}$$

Similarly, if the first qubit in the example above (the first qubit in the topmost equality in Figure 1.3) was initially in the state $|-\rangle$, rather than in the state $|+\rangle$, we would, for the outcome $s = 0$ obtain the output state $HZ_{-\phi}|-\rangle$. Since the states $|+\rangle$ and $|-\rangle$ form an orthonormal basis, for an arbitrary first "input" qubit state $|\psi\rangle$ we would, by linearity, obtain the output $HZ_{-\phi}|\psi\rangle$, and the same would hold for arbitrary general mixed states. This simple computational process easily generalizes when one considers a chain of qubits. Consider the following 5 qubit chain example as illustrated in Figure 1.3, where the first qubit is in some arbitrary input state $|\psi\rangle$ and the rest are pre-set to the $|+\rangle$ state. If the measurements performed are characterised by measurement angles $-\phi_1, -\phi_2, -\phi_3$ followed by a fixed measurement of the Pauli-X observable (*i.e.* angle 0) the resulting state, in the case all measurement report the $s_i = 0$ outcome the output state

is:

$$HHZ_{\phi_3}HZ_{\phi_2}HZ_{\phi_1}|\psi\rangle = Z_{\phi_3}\left(HZ_{\phi_2}H\right)Z_{\phi_1}|\psi\rangle = Z_{\phi_3}X_{\phi_2}Z_{\phi_1}|\psi\rangle, \tag{1.7}$$

with $X_\phi := |+\rangle\langle+| + e^{i\phi}|-\rangle\langle-| = e^{-i\phi X/2} = HZ_\phi H$. It can be easily shown that such a sequence of rotations is sufficient for the realization of any single qubit unitary [10, 34]. Intuitively, the $U(2)$ operators Z_ϕ, and X_θ rotate the Bloch sphere vector about the Z and X basis respectively, so that latter expression can be seen as an Euler decomposition of an arbitrary rotation (in $SO(3)$) rotating the Bloch vector of the input state $|\psi\rangle$. Thus, by using 4 auxilliary qubits, an arbitrary single qubit unitary can be implemented (up to the irrelevant global phase), provided all the measurements result in the $s = 0$ outcome. The "all zeroes" sequence of measurement outcomes which ensures the desired input state unitary evolution is commonly referred to as *the positive branch*. This post-selected measurement-based computation over such a qubit chain–line can then be understood as a sequence of single qubit unitaries acting on one qubit as is illustrated in Figure 1.3.

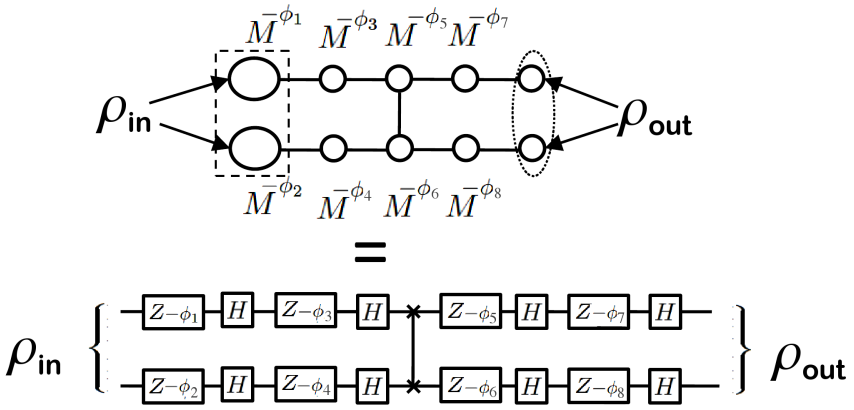

Figure 1.4. Qubit lines and entanglement in MBQC. If the resource state comprises multiple qubit lines which are entangled, single qubit unitaries and two qubit gates can be implemented by measurements. The multiple qubit chains which are entangled can be viewed as multiple qubit computational lines in the circuit model which have an entangling gate acting on them, as illustrated in the two bottom images.

The second component needed for universality computation, aside from single qubit unitary gates, is a two-qubit entangling gate. In MBQC, two-qubit gates are also realized by using pre-existing entanglement between multiple "single qubit computational lines" 1.4 .

So far we have illustrated how by using:

- particular graph states
- single qubit measurements restricted to outcomes in the XY plane of the Bloch sphere

one can implement individual single qubit gates and specific two-qubit gates in the positive branch.

To achieve computational universality the single qubit gates and two-qubit gates have to be combined into a larger construction. Such "larger constructions" can be designed in a generic way, which brings us to the concept of *generic families of graph states* In particular, we focus on a family of graph states called *brickwork states*, defined in Figure 1.5. The family of brickwork states are a universal resource for the one-way model, in the positive branch, which can be shown by using the universality of the circuit model, as given by the following lemma:

Lemma 4. *For every $n-$qubit quantum circuit of size S, which uses the Hadamard, $\frac{\pi}{8}$ and nearest-neighbour $CNOT$ gates, the corresponding unitary evolution can be implemented in the one-way model, using single qubit XY plane measurements only, using a brickwork state of n rows and (at most) $9S$ columns in the positive branch.*

Each "brickwork block" (see Figure 1.5) can implement both the Hadamard and the $\frac{\pi}{8}$ gate on any of the two computational lines. Also any "brickwork block" can implement a CNOT gate, and the identity on the two qubit lines. These two claims are shown in [1]. Then to generate any single qubit unitary (in the set we consider) we need at most one column of brickwork blocks (so 5 columns of the brickwork state). However, the brickwork blocks come in interlaced pairs, so to generate a CNOT between any two neares neighbours, we need at most 2 columns of brickwork blocks, thus at most 9 columns of the brickwork state. In total, this the entire translation of an arbitrary circuit as in the statement of the lemma will require at most $9S$ brickwork columns □.

Since a quantum circuit which allows the Hadamard, $\frac{\pi}{8}$, and CNOT gate is quantum computationally universal, so is one way computation over the brickwork state, in the positive branch.

In our initial definition of the circuit model the quantum machine performed single and two-qubit gates over an $n-$qubit register, initially pre-set to the $|0\rangle^{\otimes n}$ state. More generally, the same machine can perform any unitary evolution over any $n-$qubit state ρ, by initially pre-setting the register to the state ρ_{in}. Also, in principle, the desired output of a computation can be a quantum state ρ_{out} residing in the register after all the quantum gates have been applied. The presented model of one-way computation allows for such a quantum input and quantum output as well. This is achieved by selecting a subset of the qubits of the resource state to be the input partition, and a subset of qubits which will be the output partition, hence, unmeasured. This we have hinted in figures 1.3 and 1.4 by designating a quantum state to the leftmost qubits in both illustrations. Similarly, the rightmost qubits did not have designated measurement angles, signifying they will remain unmeasured and the output of our computation. In the case of the brickwork state, usually the leftmost column of the brickwork is reserved for the input, and the rightmost for the output.

Which partitions we choose to be reserved for the input and output influence the computational properties of the underlying graph states. This we will return to shortly. The concept which generalizes the graph of the underlying graph state, and takes into account the input and output partitions is called an *open graph state* [35]:

Definition 5. *An* open graph state *is a triplet (G, I, O) of a graph $G = (V, E)$, and two subsets of the set of vertices of the graph $I, O \subseteq V$.*

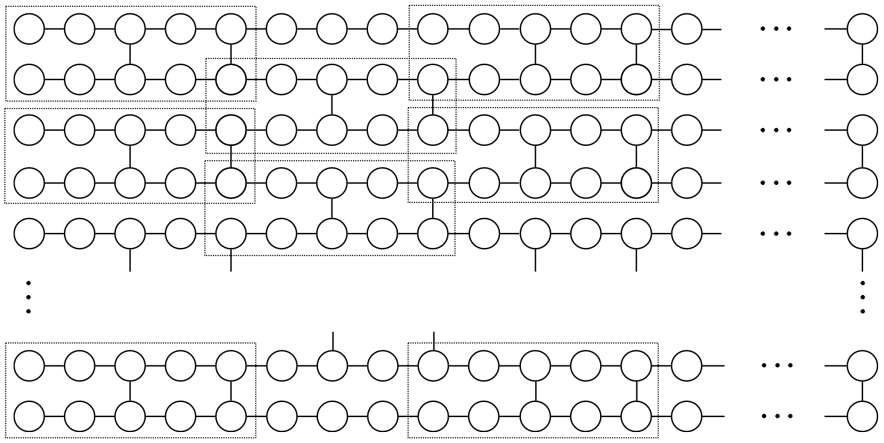

Figure 1.5. The brickwork state*, $\mathcal{G}_{n\times m}$, a universal resource state for measurement-based quantum computing requiring only single qubit measurement in the (X, Y) plane [1]. Qubits $|\psi_{x,y}\rangle$ $(x = 1, \ldots, n, y = 1, \ldots, m)$ are arranged according to layer x and row y, corresponding to the vertices in the above graph, and are originally in the $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ state. Controlled-Z gates are then performed between qubits which are joined by an edge. The rule determining which qubits are joined by an edge is as follows: 1) Neighbouring qubits of the same row are joined; 2) For each column $j = 3 \mod 8$ and each odd row i, the qubits at positions (i, j) and $(i + 1, j)$ and also on positions $(i, j + 2)$ and $(i + 1, j + 2)$ are joined; 3) For each column $j = 7 \mod 8$ and each even row i, the qubits at positions (i, j) and $(i+1, j)$ and also on positions $(i, j + 2)$ and $(i + 1, j + 2)$ are joined. The building blocks (emphasized by dotted rectangles) are each individually capable of implementing any single qubit unitary, or a two-qubit entangling gate.*

Now we can define the process of (probabilistic) one-way computation as given in Protocol 1.2.1. Which computation we perform in this model of computation depends on the geometry of the open graph state and the choice of computational angles. If the open graph state is the brickwork state where the first (leftmost) column of qubits is chosen as the set of inputs I and the last (rightmost) as O then, given a sufficiently large brickwork state, any unitary transform of the input qubits can be performed.

Every open graph G and every sequence of measurement angles $\{\phi_k\}_k$, chosen for the probabilistic process described in 1.2.1 defines a map $\Gamma_G(\{\phi_k\}_k)$. This is a linear map (more precisely, a completely-positive map) from the Hibert space of the input qubit partition to the Hilbert space of the output qubit partition. To understand what a particular sequence of angles over a given graph realizes, in other words, what $\Gamma_G(\{\phi_k\}_k)$ is, one can translate the one-way quantum computation directly into the circuit model. Many techniques have been developed for translations between the two (equivalent) models. By translation and the relationship between a circuit and the map it implements one can understand what a particular one-way computation does. However, it is possible to more directly address the properties of the map $\Gamma_G(\{\phi_k\}_k)$ and we have done so in an analysis we call the Generalized Phase Map Decomposition, presented in Section 5. As a side result of this analysis, we capture some of the known results on classical simulability of MBQC using somewhat different approaches.

Protocol 2 Probabilistic one-way computation

1. For a chosen input $n-$qubit state ρ_{in}, an open graph state $(G = (V, E), I, O)$ is selected with $N = |V|$ vertices and n vertices in the input partition, so $n = |I|$. The qubits in ρ_{in} are labelled $1, \ldots, n$, the vertices V are integer labels $1, \ldots, N$ such that the vertices in I are first n labels, and the output vertices (in O) are the last $m = |O|$ labels.
2. $N - n$ qubits pre-set in the state $|+\rangle$, labelled $n + 1, \ldots, N$ are added to the state ρ_{in}.
3. The entangling $\wedge Z$ operation is applied to all the qubits whose labels are neighbours in the graph G.
4. A sequence of *computational measurement angles* $(\phi_1, \ldots, \phi_{N-m})$ is chosen.
5. The measurements M^{ϕ_i} is applied on the qubit with label i, for all $i \in V \setminus O$.
6. If all measurements collapse to the positive branch, the computation is successful, and the remaining quantum state of unmeasured $m-$qubits is the desired output. .

The presented model of computation is probabilistic: the map we desire is realized only if all the measurements collapse to the positive branch, the probability of which in general on the order of $1/2^S$, where S is the number of qubits we measure. The key property of MBQC in general (and the one-way model in particular) is that, provided the underlying resource has certain properties, the positive computational branch can be deterministically simulated by imposing an *adaptive measurement strategy*. That is, by fixing the order of measurements, and by allowing the measurement angles which define the computation in the positive branch to depend on the prior measurement outcomes. This we explain presently, and show how *determinism* in the one-way model can be achieved by adaptive measurement strategies.

1.2.2 Determinism in the one-way model

We begin by illustrating how determinism can be achieved in the simplest non-trivial case: a single qubit measurement over a 2 qubit chain, as illustrated in Figure 1.3.

As we have shown, provided the input state (the state of the first qubit) was $|+\rangle$, and the measurement was defined by the angle ϕ, in the case the measurement outcome s was 0, the output state is $HZ_{-\phi}|+\rangle$. Recall, this is derived by applying the (projector) $(\langle+|Z_{-\phi} \otimes \mathbb{1})$ on the joint post-entanglement state of the two qubits. For the case of the opposite outcome, we should apply the "negative branch" projector: $(\langle-|Z_{-\phi} \otimes \mathbb{1}) = (\langle+|Z_{-\phi}Z \otimes \mathbb{1})$, which is equivalent to phase-flipping (using a Pauli-Z operation) the first qubit and subsequently applying the positive branch projector. This has the following implication: the output state of a two qubit one-way computation in the case of the unwanted outcome $s = 1$ would be equal to the state in the desired positive branch (s=0) if we could "anachronically" apply a Pauli-Z correction on the pre-measurement graph state. This is clearly impossible.

However, note that it is possible to apply a Pauli-Z operation on the first qubit, without disturbing the total two-qubit graph state - one simply needs to additionally apply a Pauli-X gate on the second qubit as well. This observation comes from the theory of *stabilizer states* [32, 10].

To define the stabilizer states we first define the $n-$qubit Pauli group:

Definition 6. *The $n-$qubit Pauli group is defined as follows:*

$$\mathcal{P}_n = \{\pm 1, \pm i\} \times \{\bigotimes_{k=1}^{n} P_k | P_k \in \{I, X, Y, Z\}\},$$

and contains all combinations of single qubit Pauli operators which act on an $n-$qubit system, additionally equipped with a global phase from $\{\pm 1, \pm i\}$. This global phase is required to make this set a multiplicative group.

A set S of n independent operators in $\mathcal{P}_n$ has a unique state $|\psi\rangle$ such that for each $P \in S$ the state $|\psi\rangle$ is the $+1$ eigenvalue eigenstate of P. The set S, and more generally the subgroup $\langle S \rangle$ is said to *stabilize* the state $|\psi\rangle$, which is then called a *stabilizer state*. All the elements in $\langle S \rangle$ are called *stabilizers*.

The mathematical theory of stabilizers is all but ubiquitous in quantum information processing. It had a profound impact on the theories of quantum error correction and fault-tolerant quantum computation [36, 10], was used in the fabled Gottesman-Knill Theorem [10] (proving Clifford computations are classically simulatable), many quantum secret sharing schemes [37], results on quantum multi-party computation protocols [38], and so on. Here, we use it to ensure determinism in the one-way model. Graph states we discussed previously are stabilizer states, and characterised as follows:

Lemma 7. *Let $|G\rangle$ be a graph state characterised by the graph $G = (V, E)$ on n vertices in the set V. Then $|G\rangle$ is stabilized by the operator*

$$K_v = X_v \prod_{w \in N(v)} Z_w, \ \forall v \in V \tag{1.8}$$

where X_v is the Pauli X operator acting on the qubit corresponding to the vertex v, Z_w is the Pauli Z operator acting on the qubit corresponding to the vertex w, and $N(v)$ is the set of all vertices which are neighbours of the vertex v in G. The state $|G\rangle$ is the unique state stabilized by the set of independent operators $\{K_v\}_{v \in V}$.

The operators K_v above are often called *correlation operators* [32].

In our case of a two-qubit connected simple graph the stabilizer generators are: $X \otimes Z$ and $Z \otimes X$. We are interested in the latter stabilizer $Z \otimes X$. Since this operator acts as an effective identity on the two-qubit graph state $\wedge Z|+\rangle|+\rangle$, and since the Pauli-Z and Pauli-X operators are both unitary and Hermitian, hence self-inverse, the following equality holds (up to irrelevant global phase):

$$(Z \otimes \mathbb{1}) \wedge Z|+\rangle|+\rangle = (\mathbb{1} \otimes X) \wedge Z|+\rangle|+\rangle. \tag{1.9}$$

Note that the right-hand side state in the equation above commutes with the measurement of the first qubit, thus instead of performing the physically impossible anachornical Z correction on the first qubit, we can simply apply a (post-measurement) X correction on the second. Since the operator $Z \otimes \mathbb{1}$ commutes with the stabilizer $Z \otimes X$, and the entangling gate $\wedge Z$ we get

the same result in the case the first qubit was in initially in the state $|-\rangle$. But then, by linearity, this post-measurement correction process works for any input state (of the first qubit.) Thus, we get the following deterministic process, which includes a controlled correction, conditional on the measurement outcome illustrated in Figure 1.6. The idea of using stabilizers to ensure determinism is generalized to our 5 qubit 1WM example we used to simulate and arbitrary single qubit unitary. The stabilizer generators of a 5 qubit 1D chain graph states are: $XZIII$, $ZXZII$, $IZXZI$, $IIZXZ$ and $IIIXZ$, omiting the symbol for the tensor product. Then, to correct for the undesired outcome of the first qubit measurement $s_1 = 1$ we use the second generator (which simulates the anacronical Pauli-Z correction), and apply an X gate to the non-measured nearest neighbour (qubit 2) and a Z correction on the distance 2 non-measured neighbour (qubit 3). In general, the undesired measurement outcome $s_k = 1$ of the k^{th} qubit is corrected by using the stabilizer with the Pauli-X operator acting on the $k + 1^{st}$ qubit.

The required corrections need not be implemented physically on the qubits which will subsequently be measured. The effect of physical corrections can be simulated by *adapting the performed measurement angle*. In particular, if the "positive branch" measurement angle ϕ is to be applied on a qubit which inherits an X correction due to the previous outcomes of measurements, the angle ϕ attains a minus sign: $\phi' = (-1)\phi$. Alternatively, if the required correction is Z then a π phase is added: $\phi' = \phi + \pi$. Both corrections may occur simultaneously, in which case simply both the sign and the phase are modified (by adding a π). The X and/or Z corrections need to physically be applied only on the output qubits. In the case they will be measured as well (in the case of a classical output of the computation), the corrections can again be done on the classical output, rather than the quantum state. This type of a correction strategy cannot be generalized to all computation-underlying graph states. In more complicated graphs this strategy would require corrections to be applied on qubits which have already been measured which is not possible. A significant body of work exists which characterises the possible correction strategies, and the properties of the underlying resource states which are required to ensure determinism by adaptive strategies. In particular, the concepts of *flow* and *generalized flow* (gflow) [35, 39] give graph-theoretical characterisations of the geometry of the underlying resource state required to ensure determinism and provide the adaptive corrective strategies. As we have stated in Protocol 1.2.1, a (probabilistic) one-way computation is defined by an open graph state, and a sequence of measurement angles. Whether a particular open graph state allows for deterministic computation, depends on the graph-theoretical property of flow, which we now define.

Definition 8. *An open graph state has* flow *if there exists a map $f : V \setminus O \to V \setminus I$ (from measured to non-input qubits) and a partial order $\succ$ over V such that for all $i \in V \setminus O$:*

F1 The edge $(i, f(i))$ is in the set of edges E,

F2 $f(i) \succ i$, and

F3 for all neighbours k of $f(i)$ except i, we have $k \succ i$.

The property of flow has an operational consequence: if an open graph state has flow, then deterministic computation (*i.e.* simulation of the positive branch of computation, as discussed previously) can be implemented using an adaptive measurement strategy. Moreover, the flow

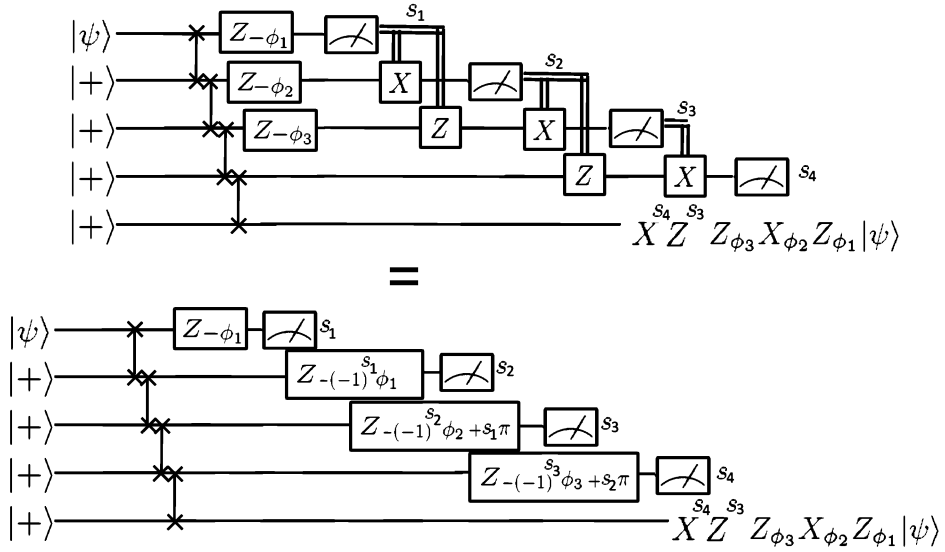

Figure 1.6. Correction strategies in MBQC. In MBQC, non-physical projection (which we have seen is sufficient for universal quantum computation over suitable resource states) can be simulated by adaptive strategies. Undesired measurement outcomes are compensated for by local stabilizer operations, conditional on the measurement outcomes, as explained in the main text. In the bottom image the same effect is achieved not by directly applying unitary X and Z corrections, but rather by modifying the future measurement angles. The "circuit" in the bottom image is non-standard as the gates applied depend on measurement outcomes of other qubits. This exemplifies the difference between MBQC and the circuit model.

construction dictates how the adaptive measurements have to be performed, and this is the main content of theorem 1 in [35].

The partial order $\succ$ introduces the order in which qubits have to be measured, and partitions the set V into disjoint subsets $V_1, \dots V_l \subseteq V$, such that for all subset labels k, j, $k > j$ implies $w \succ v$ for all $w \in V_k$ and for all $v \in V_j$. The subsets V_j are called *layers*. The flow function f is bijective, and determines how measurement angles need to be adapted, based on previous measurement outcomes. To each qubit/vertex v in layer V_k we can assign a set of vertices D_v^X and D_v^Z such that all the vertices in D_v^X and D_v^Z are in layers preceding V_k, with respect to the order $\succ$. A qubit/vertex v is in D_w^X for a vertex w if $f(v) = w$. A qubit/vertex v is in D_w^Z for a vertex w if $v \in N(f(v)) \setminus \{v\}$, where with $N(v)$ we denote the neighbourhood of v in G.

The set D_w^X is called the *set of X dependancies of the qubit w* and the set D_w^Z is called the *set of Z dependancies of the qubit w*. Qubits in the first layer have no dependencies, and qubits in the second layer have no Z dependencies.

Now we can give a single party protocol, which shows how deterministic computation can be

done over open graph state with flow. Assume the desired computation is defined by an open graph state $(G = (V, E), O, I)$ with flow $(f, \succ)$, and the sequence of measurement angles $(\phi_1, \ldots, \phi_{N-m})$, as given in the Protocol 1.2.1. Labelling is performed as in the probabilistic protocol. Then, the deterministic version of this computation is given with Protocol 3.

Protocol 3 Deterministic one-way computation

1. $N - n$ qubits pre-set in the state $|+\rangle$, labelled $n + 1, \ldots, N$ are added to the $n-$qubit input state ρ_{in}.
2. The entangling $\wedge Z$ operation is applied to all the qubits whose labels are neighbours in the graph G.
3. The layers $V_1, \ldots V_l$, respecting the order $\succ$, along with dependency sets D_i^X and D_i^Z are determined for each qubit.
4. For each layer label $x = 1 \ldots l - 1$
   For each label $v \in V_x$
   measure the qubit labeled v with an *adapted measurement angle* ϕ'_v and store the measurement outcome $s_v \in \{0, 1\}$ called the signal. The adapted measurement angle depends on the corresponding *computational* angle (ϕ_v) and prior signals as follows:
   $$\phi'_v = (-1)^{\oplus_{j \in D_v^X} s_j} \phi_v + (\bigoplus_{j \in D_v^Z} s_j)\pi, \tag{1.10}$$
   where $\oplus$ denotes modulo 2 addition.
5. For each label $v \in V_l$ (output qubits)
   if $\bigoplus_{j \in D_v^X} s_j = 1$ apply an X gate to qubit labeled v
   if $\bigoplus_{j \in D_v^Z} s_j = 1$ apply a Z gate to qubit labeled v. .

The flow and gflow constructions which enable deterministic computation induce a temporal order on the measurements. This defines a concept of *computational depth* in MBQC, in terms of the number of layers we require. It was shown that, if classical processing is ignored, the MBQC model may offer computational speed-up with respect to computational depth when compared to the circuit model. For an overview of these topics we refer the interested reader to the Master's thesis [40] and also to [41].

The one-way allows for an arbitrary unitary evolution over suitable generic resource states, assuming the capability of a continuum of possible measurement settings. In practice, however, this is impossible. But, just as in QCM the $\frac{\pi}{8}$ gate, a Hadamard gate and a two-qubit CNOT gate are sufficient for *approximate* universality, in the one-way model it will suffice to restrict the measurement angles to the finite set of 8 equidistant angles of the form $\left\{\frac{k\pi}{4}\right\}_{k=0}^{7}$. A restriction to Pauli bases measurements only would suffice for any arbitrary Clifford computation.

Throughout the rest of Part 1 of this thesis, all the angles which appear are always members of this approximately universal discrete set, although the constructions we present would work with a continuum of masurement settings as well.

1.2.3 UBQC from the one-way model

A deterministic one-way quantum computation conceptually comprises two parts: a *classical control unit* and a *quantum unit*. The quantum unit sets up the scene - generates the relevant (generic) computational resource state - and is capable of performing single qubit measurements of the observable M^ϕ (XY Bloch sphere plane measurements) on a desired qubit. The classical unit directs the measurement angles to the quantum unit, which are when needed adapted, based on the measurement outcomes reported by the quantum unit.

The central idea behind UBQC is to use this separation and allocate the classical controller unit to the client and the quantum unit to the server. To ensure privacy, however, the computation needs to be encoded: this is achieved in UBQC by effectively encoding the resource state.

The standard procedure in the one-way model, to prepare the resource state, is to start with a set of qubits in a fixed state, say $|+\rangle := |+_0\rangle$, and to apply an entangling operation of the controlled phase gate ($\wedge Z$) to some of them.

In UBQC, in contrast, the client will provide the initial phase rotated qubits of the form $|+_\theta\rangle$ to the server, without informing him of the values of $\theta \in \{0, \pi/4, \ldots, 7\pi/4\}$. Applying the entangling gates then prepares an encoded resource state [5]. Now, if one was to measure a qubit in the usual MBQC protocol with some measurement angle ϕ, this would be equivalent to measuring the pre-rotated qubit (in the state $|+_\theta\rangle$) with the angle $\delta' = \phi + \theta \mod 2\pi$, as the phase rotation and $\wedge Z$ gate commute. In this case, the measurement angle alone says nothing about the computation run, but a malicious server may still learn something about θ when given δ', hence also about ϕ (*i.e.* about the computation). To see this, note that the server knows the relationship $\delta' = \phi + \theta \mod 2\pi$, (as we must assume the server knows the protocol). Once the server is given δ' he can apply the rotation $Z_{-\delta}$ on the qubit $|+_\theta\rangle$ obtaining the qubit in the state $|+_{-\phi}\rangle$. By applying an X gate he obtains a qubit in the state $|+_\phi\rangle$, from which, by measurement, he can learn (up to) 1 bit of information about the angle ϕ.

To solve this security loophole, UBQC exploits the probabilistic nature of MBQC. The client sends a modified measurement angle $\delta = \phi + \theta + r\pi \mod 2\pi$, where $r \in \{0, 1\}$ is chosen randomly by the client and hidden from the server. The value of r can be interpreted as a flip of the measurement outcome, which can be easily compensated by the client. To gain intuition why the additional bit of randomness conceals all of the secret information, note that if the server in this case proceeded again by applying the $Z_{-\delta}$ rotation followed by an X gate, the qubit he obtains is in the state $1/2|+_\phi\rangle\langle+_\phi| + 1/2|+_{\phi+\pi}\rangle\langle+_{\phi+\pi}|$ (corresponding to the choices $r = 0$ and $r = 1$ of the random parameter r). This is a maximally mixed state, so no longer correlated to ϕ. This argument we make more formal presently.

Now the quantum information (pre-rotated qubits) and classical information (measurement angles) accessible to the server is no longer correlated to the client's desired computational angles

[5] The pre-rotation of the qubits, performed using the Z_θ gates commutes with the entangling operation $\wedge Z$, so the correctness of the computation can still be maintained. This is a key property of the resource state in this model of computation which allows this type of encryption to work. We will address this property to more detail in Section 3.3.1.

(denoted ϕ), and this constitutes the crux of the proof of blindness of UBQC [1]. A quick summary of the UBQC protocol is given next, and a full detailed breakdown is given in Protocol 4.

Initially, in the *preparation phase*, the client sends S (the *size* of the computation) randomly pre-rotated qubits in the states $\{|+_{\theta_i}\rangle\}_{i=1}^{S}$, to the server, keeping the angles θ_i secret. The server then builds up the brickwork state using the received qubits and the $\wedge Z$ interaction. Proceeding sequentially on each qubit, if the desired measurement angle for qubit i was ϕ_i (defined for the non-prerotated resource state, and including the necessary adaptations to the angle based on prior measurement outcomes $s_{k<i}$), the client will ask the server to measure the qubit with respect to the angle $\delta_i = \phi_i + \theta_i + r_i\pi \mod 2\pi$, where the binary parameter r_i is chosen randomly. The server reports each measurement outcome s_i which the client flips if $r_i = 1$.

In the case of an honest server, this procedure yields the correct outcome of the computation. Moreover, regardless of the malicious activity of the server the client's privacy is unconditional - the protocol is blind. Before we can proceed with giving a proof sketch of blindness as given in the original paper [1], we need to formalize the concept of blindness a bit, and explicitly give the actual protocol. For simplicity, we shall focus on the classical input - classical output setting meaning the client includes the preparation of his actual desired input state in the overall computation, and that the final quantum output is measured out by the server.

Suppose the client has in mind a unitary operator U that is implemented with a particular sequence of measurements on a brickwork state $\mathcal{G}_{n\times m}$ (Figure 1.5) with measurements given as multiples of $\pi/4$ in the (X, Y) plane with overall computation size $S = n \times m$. This pattern could have been designed either directly in MBQC or from a circuit construction. Each qubit $|\psi_{x,y}\rangle \in \mathcal{G}_{n\times m}$ is

indexed by a *column* $x \in \{1, \ldots, n\}$ and a *row* $y \in \{1, \ldots, m\}$. Thus each qubit is assigned a measurement angle $\phi_{x,y}$, a set of X-dependencies $D_{x,y} \subseteq [x-1] \times [m]$ and a set of Z-dependencies $D'_{x,y} \subseteq [x-1] \times [m]$. Here, we assume that the dependency sets $X_{x,y}$ and $Z_{x,y}$ are obtained via the flow construction.

During the execution of the pattern, the actual measurement angle $\phi'_{x,y}$ is computed from $\phi_{x,y}$ and the previous measurement outcomes in the following way: let $s^X_{x,y} = \oplus_{i\in D_{x,y}} s_i$ be the parity of all measurement outcomes for qubits in $X_{x,y}$ and similarly, $s^Z_{x,y} = \oplus_{i\in D'_{x,y}} s_i$ be the parity of all measurement outcomes for qubits in $Z_{x,y}$. For the special case of $x = 1$ we define $s^X_{1,y} = s^Z_{1,y} = 0$. Then $\phi'_{x,y} = (-1)^{s^X_{x,y}}\phi_{x,y} + s^Z_{x,y}\pi$.

We assume that the client's input and output of the computation are built into U. In other words, the client wishes to compute the results of some fixed single qubit measurements in the (X, Y) plane of the state $U(|+\rangle \ldots |+\rangle)$. This protocol easily extends to deal with arbitrary classical or quantum input and output [1].

In this case, the input state the client prepares and forwards to the server needs to be quantum one-time padded by the application of random Pauli-X gates and relative phase randomized Z_θ gates.

However, throughout this thesis we will mostly be dealing with classical input and classical output scenarios, even though almost every claim we will make can easily be adapted to work with a quantum input as well.

Protocol 4 implements a blind quantum computation for U.

Protocol 4 Universal Blind Quantum Computation

1. **Client's preparation**
   For each column $x = 1, \ldots, n$,
   for each row $y = 1, \ldots, m$,
   (a) the client prepares the state $|\psi_{x,y}\rangle \in \{|+_{\theta_{x,y}}\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta_{x,y}}|1\rangle) \mid \theta_{x,y} = 0, \pi/4, \ldots, 7\pi/4\}$, where the defining angle $\theta_{x,y}$ is chosen uniformly at random, and sends the qubits to the server.

2. **Server's preparation**
   (a) The server creates an entangled state from all received qubits, according to their indices, by applying CTRL-Z operators between the qubits in order to create a brickwork state $\mathcal{G}_{n\times m}$.

3. **Interaction and measurement**
   For each column $x = 1, \ldots, n$
   For each row $y = 1, \ldots, m$
   (a) The client computes $\phi'_{x,y}$.
   (b) The client chooses a binary digit $r_{x,y} \in \{0, 1\}$ uniformly at random, and computes $\delta_{x,y} = \phi'_{x,y} + \theta_{x,y} + \pi r_{x,y}$.
   (c) The client transmits $\delta_{x,y}$ to the server, who performs a measurement in the basis $\{|+_{\delta_{x,y}}\rangle, |-_{\delta_{x,y}}\rangle\}$.
   (d) The server transmits the result $s_{x,y} \in \{0, 1\}$ to the client.
   (e) If $r_{x,y} = 1$, the client flips $s_{x,y}$; otherwise she does nothing.

As we have explained, the brickwork state above is a generic and universal resource for one-way quantum computation, where the single qubit measurement which drive the computation are restricted to the (X, Y) plane of the Bloch sphere. The generic nature of this resource guarantees that no information about the computation can leak to the server on the basis of the geometry of the underlying resource.

Throughout the rest of this thesis, for simplicity, we will use one-dimensional indexing for all the involved parameters: the measurements angles ϕ_i, the random variables r_i and θ_i, the messages to the server δ_i (the measurement angles to be performed) and the server's messages to the client s_i (measurement outcomes), for $i \in \{0, \ldots, N-1 = n \times m - 1\}$. In particular, one has $\delta_i = (-1)^{s_i^X}\phi_i + s_i^Z\pi + \theta_i + r_i\pi$. It was shown in [1] that this protocol is correct, *i.e.*, if both the client and the server follow the steps of the protocol then the classical outcome are the results of some fixed single qubit measurements of the state $U(|+\ldots+\rangle)$. These classical outcome corresponds to the signals s_i generated by the measurements of the final layer of the brickwork state then bit flipped if the corresponding parameter r_i was 1 and left as they are if the corresponding r_i was zero. Moreover, the server will not learn anything about the client's computation, *i.e.*, the protocol is *blind* with respect to the following definition:

Definition 9. *We say a Protocol P on input X is* blind while leaking at most L(X)*, where $L(X)$ is any function of the input if:*

1. *The distribution of the classical information obtained by the server in P depends only on $L(X)$.*
2. *Given the distribution of classical information described in 1. and $L(X)$, the state of the quantum system obtained by the server in P is fixed.*

This definition, used in [1, 42] is effectively a generalization of the definition suggested in [23] (see Definition 2) to the quantum information processing realm. This definition gives a family of characterizations of blindness where differing types and amounts of information are allowed to be leaked to the server. For example, as we will see, in the case of UBQC, the input X from the definition comprises the input to the computation the client wishes to perform, and the description of the computation–the computational angles $\phi_{x,y}$. The function L, known as *the leakage function*, for the case of UBQC, upon the input X simply returns the size of the computation, *i.e.* the dimensions $n \times m$ of the resource state. What this definition then says is that the classical information the server gets is independent from every aspect of the input X, except from the size (as the server clearly has to be told how to build the brickwork state, its dimensions), and that all the quantum systems the server obtains from the client, even when all the classical information he obtained is taken into account, is also independent from any aspect of X except the size. As the upper bound of the size can never be concealed [6] this is the best that can be done in principle. The function L could characterize other types of blindness characteristics as well, where, say, the computation is leaked (meaning, $L(X)$ returns the computational angles from X, but not the input), or like in the case of the protocol of Childs (see Section 1.1), where some aspects of the computation are leaked but not all of it [7].

Theorem 4. *[1]* **Protocol 4** *is blind while leaking at most the dimensions of the brickwork state, i.e. an upper bound on the input size and the depth of the computation.*

Strictly speaking, in UBQC, the server learns additional information coming from the structure of the protocol. In particular, given a UBQC computation was run, of size $n \times m$ the server learns that the client is implementing a unitary which *can* be implemented on a brickwork state of that size using the eighth measurements allowed. The upper bound on size of the computation can never realistically be fully concealed, if the server knows when the computation has commenced and when it terminated, which will surely be the case in any sensible delegated computation scheme, as we mentioned in Footnote 6. Thus, in UBQC one focuses on what can be concealed - in this case, it is the computational angles (ϕ_i) which define the clients computation.

The main idea of the proof of blindness, *i.e.* of the Theorem 4 can be summarized as follows.

Proof idea:

[6] Aside from trivial protocols, where the client computes the computation herself and sends irrelevant information or nothing to the server, if the server actually *does perform the computation*, he knows when it began and ended, and from this he can always infer the upper bound on the size.

[7] In the protocol of Childs, the server learns when the client wishes non-Pauli single qubit gates or the CNOT to be performed, but not the Pauli gates, as these can be performed by the client herself.

The criterion 1 of the definition of blindness is almost trivially satisfied: for the i^{th} qubit, the server receives the angle: $\delta_i = \phi'_i + \theta_i + r_i\pi$. Since θ is uncorrelated with the computational angles, and the sum is taken modulo 2π, it effectively one-time pads the angle ϕ'_i (the adapted angle) which is the only component of δ_i correlated to the computational angles. Thus the classical information is uncorrelated to the input of the client.

Concerning the quantum information, the criterion 2 of blindness, the server is, for the i^{th} qubit in possession of two types of information: a qubit in the state $|+_{\theta_i}\rangle$ and the measurement angle $\delta_i = \phi'_i + \theta_i + r_i\pi$. The state of the qubit can then be rewritten as:

$$|+_{\theta_i}\rangle\langle+_{\theta_i}| = |+_{\delta_i-\phi'_i+r_i\pi}\rangle\langle+_{\delta_i-\phi'_i+r_i\pi}| = Z^r|+_{\delta_i-\phi'_i}\rangle\langle+_{\delta_i-\phi'_i}|Z^r. \tag{1.11}$$

Now, although the server knows δ, for every fixed ϕ'_i the state above is in one of two orthogonal qubit states depending on r_i. Since this parameter was chosen uniformly at random, at this particular step of the computation, the server has no access to it. Thus, his system is in a equiprobable mixture of orthonormal qubit states. This is a totally mixed state, for any choice of computational angles by the client, hence the second criterion of blindness holds as well. □

However, a slightly different proof of blindness, used in our formalism of approximate blindness, will be given in Chapter 2. In Chapter 4, Section 4.3, a new, perhaps more detailed proof of blindness of the original protocol is presented.

1.2.4 Variants of the UBQC protocol and the two server setting

In the presentation of the UBQC protocol we have focused on the "classical input" variant of the protocol in which the concealed computation also encodes the input of the client. However it is possible for UBQC to work with a fully quantum input from the client, in which the quantum input is also fully concealed from the server. This is achievable in two very related ways. For illustration purposes, let $|\phi\rangle$ be a single qubit input state the client wishes to use as her input. The client's input could in general be any $n-$qubit mixed or pure state. To conceal her input, the client applies a quantum one-time pad [43] on her state, thus sends the state $X^xZ^z|\phi\rangle$ to the server, where the binary parameters x and x are chosen uniformly at random. This state reveals no information about the input to the server. The client sends this state, along with all the required pre-rotated qubits for the computation to the server. The server assembles the pre-rotated resource state, and teleports the state $X^xZ^z|\phi\rangle$ onto the (known) input qubit. To achieve the teleportation the server will simply prepend the input qubits to the rest of the resource state as the first layer, by using the $\wedge Z$ interaction. The first layer is then measured with the fixed Pauli-X measurement. This teleports the one-time input onto the pre-rotated qubits up to a local Hadamard gate H. From here the computation proceeds as we have described for the classical input case, and the client adapts the measurement angles to counteract the quantum one-time pad, and the undesired Hadamard gate (which will require the client to send one extra computational qubit) in run time.

Alternatively, the client applies a more involved one-time pad to her input quantum state of the

form $X^x Z_{\theta_0}|\phi\rangle$, in which case the input state is directly used to build up the encrypted resource state. Her measurement angles are again adapted to ensure correctness. Both variants can be easily shown to ensure complete privacy for the client, for details see [1]. Both approaches do require that the client either has a one-time padded quantum state to begin with, or that she has the power to apply the quantum one-time pad herself. For multi-qubit inputs, the same procedure is applied on each of the input qubits.

Similarly, the outputs of UBQC can be quantum (in which case, the Pauli-X and Z byproducts on the output state, which are concealed from the server offer a quantum one-time pad of the output state, and the client will need to undo those herself), or classical. In the classical case the server measures out the one-time padded quantum output with respect to a fixed observable, say Pauli-Z. In this case, the classical measurement output is classical one-time padded with respect to the hidden Pauli-Z byproducts of the quantum output state.

The two server setting In the UBQC protocol, the client is assumed to have modest quantum powers: she can prepare and emit specific single qubit states. An obvious question is: could the client be completely classical, and yet be guaranteed perfect privacy?

Already in the original proposal [1], the authors give a setting where this is possible. If one considers *two* quantum servers, who share a sufficient number of Bell pairs, but are *not allowed to communicate*, and share private channels with the client, then the client can be completely classical. The idea is simple. The client uses one server to generate the pre-rotated qubits, according to her random angles. Then, the first server teleports those qubits to the second server, and informs the client of the classical outcomes of the teleportation procedure. The client then just runs the single server protocol with the second server. Because the two servers cannot communicate (despite the shared Bell pairs due to no signalling) the clients perfect privacy is easily reduced to the blindness of the single server protocol.

While the no-communication assumption may seem curious, this setting appears in classical and quantum complexity classes of *multiprover interactive proofs systems*. We will return to the question of feasibility of a fully classical client, and what it could entail, the impact of two server settings on complexity theory and other related questions in Chapter 4. The two server setting will be revisited in Chapter 3, where we will also briefly address the question whether the no communication assumption could ever be justified in practice. For the time being, we focus our attention on the practical aspects of the single server UBQC.

In order for UBQC to make the transition from theory to practice, many issues need to be resolved. Since this is a two-party protocol, there are two sides of the story: the client's, whom we shall call "Alice", and the server's, whom almost by necessity, we call "Bob". Now we address the issues our two players may encounter when the protocol attempts the transition from ideal settings to more realistic ones.

1.3 Realizability of UBQC

1.3.1 The world according to Alice

The client Alice wishes to compute a quantum computation with the help of a friendly neighbourhood quantum server Bob while maintaining her privacy, and minimizing her required resources. If we allow Alice minimal quantum powers, minimal classical computing powers, and classical communication, then in the ideal world her problem will be solved with UBQC.

As we have seen, in the UBQC protocol correctness of Alice's computation (when Bob is honest) and her privacy (for any malicious activity of Bob) is guaranteed, provided Alice can generate perfect single qubit states, and send them off to Bob. While from a theoretical point of view, this may constitute the lowest possible quantum requirement, from a pragmatic point of view, generation of such states to be sent along long distances can never be achieved perfectly. In reality, we need to ensure that correctness and blindness (privacy) can be maintained even in presence of realistic imperfections.

Even if Alice's preparation stage, and Bob's computer are imperfect, the correctness of Alice's computation can be guaranteed if the (blind) computation Alice runs is embedded in a fault tolerant code (see [1]). However, the *blindness* of the UBQC protocol has only been established in the ideal case where the client prepares perfect qubits. In any physical implementation, however, the preparation will inevitably be imperfect and this has to be taken into account before making any statement about blindness.

To illustrate the problem let us imagine one of the simplest possible realizations of Alice in real physical systems. Qubits are easily encoded in the polarization of a single photon generated by a realistic single photon source. Such encoding can be done with high precision. Losses in transmission cause no problems for Alice – if a photon is lost during transmission in the preparation stage, Alice can simply send another one, with a new random polarization, without jeopardizing blindness. This assumes heralding on the side of the server, *i.e.* the server has to detect whether he received a photon or not, and report to Alice. Note that this is not a security jeopardizing issue, since the polarizations between the qubits are not correlated. Noise can also arguably be managed provided it does not exceed the thresholds of the fault tolerant computation Bob will run.

The problem that does arise for Alice has also been an issue in single photon (and later weak coherent pulses) schemes for quantum key distribution: in practice, completely suppressing the probability of inadvertently sending two or more identically polarized photons instead of one is very difficult, yet such an event would invalidate the privacy guarantees for the client[8].

Recently, a small-scale realization of UBQC was implemented using single-photon polarization encoding [44]. Since this was a first ever demonstration of the protocol, the security of the scheme was addressed in a post-selected setting on the emitted photon numbers. However, in a real life

[8]The classes of attacks in QKD exploiting the multi-photon emissions are often called "beamsplitting attacks".

application, perfect blindness cannot be claimed when multi-photon emissions can occur.

While the future may bring scalable and fault-tolerant quantum computation required for the server, *perfect* quantum devices required to guarantee the *perfect* security for the client are unlikely to ever be achieved in practice. This crucial observation has instigated a whole important subfield in the case of QKD, that of security of QKD under imperfections, and many issues still exist [45, 46]. For this purpose a framework of *approximate* blindness is required. However, simply quantifying how secure a UBQC protocol is, given fixed imperfections on the side of Alice (note, imperfections on the side of Bob can only jeopardize correctness, but not blindness, even if they are intentional since UBQC is unconditionally blind) is not fully satisfying. What we would ideally need is a protocol which can achieve arbitrary levels of privacy for a fixed and realistic set-up on Alice's side. These questions we address and resolve in the following chapter, and now turn our attention to even greater problems the server Bob will succumb to in the transition to realistic settings.

1.3.2 The world according to Bob

In order for UBQC to become a reality, we first have to complete a rather formidable task: we have to build Bob a quantum computer! By design, the UBQC protocol assumes the server is running, or is able to simulate, a measurement-based quantum computation over a qubit graph state of an adequate size. One of the main challenges of scalable quantum computers is the protection of the computational system from decoherence caused by the interaction with the environment. To combat this, many techniques have been developed, which include various fault tolerant computation schemes, and alternative models of quantum computation. An example of the latter are topological quantum computation proposals which keep the state of the quantum computer protected from noise by topological properties of two-dimensional anyonic physics [9].

In the recent active interaction between condensed matter physics and quantum information science, plenty of novel resource states for the measurement-based quantum computation beyond the cluster state have been proposed [48, 49, 50, 51, 52]. These new resource states have several interesting features and advantages over the graph states used in UBQC. For example, some of those resource states are gapped ground states of their parent Hamiltonians, and therefore, they can be prepared by cooling condensed matter systems and the measurement-based quantum computation can be protected from noise by an energy gap. Such protection is not possible for the case of graph state computation. More precisely, it is impossible to have a universal resource of spin-1/2 particles that is the unique ground state of a frustration-free Hamiltonian with only two particle interactions [52, 53].

In Chapter 3 we attempt to make life easier for Bob, by investigating whether a UBQC-type protocol can be designed which relies on these novel resource states for generalized MBQC. From

[9] The required conjectural non-abelian anyonic physics can also be simulated in other more standard physical systems. This approach gave rise to topological error-correcting codes and topological fault-tolerant schemes [47].

a theoretical perspective, the results of this investigation reveal an intricate interplay between classical communication, central concepts in condensed matter physics and cryptography.

Chapter 2

The world according to Alice: *UBQC under imperfections*

We introduce a quantified notion of approximate $\epsilon-$blindness, and give a generic preparation protocol which ensures arbitrary levels of blindness for a client with a fixed device such as an attenuated laser source or a realistic single-photon source

2.1 Security under imperfections

The Universal Blind Quantum Computation protocol guarantees perfectly private delegated computation for the client, provided she can prepare and emit perfect single qubit states to the server. While this may be a very modest quantum requirement, pragmatically speaking, the generation of such states to be sent along long distances can never be achieved perfectly. The issues of concern which arises from realistic imperfections are two-fold. Since the client's computation will be performed using the resource state comprising the states generated by the client herself, the correctness of the output may be jeopardized, even when the server is honest. This problem can be resolved by embedding the computation the client wishes to run in an error-correcting code. This approach has been proposed in the original paper, and recently adapted to use the Raussendorf-Goyal-Harrington [47] topological code, with an estimation of the fault tolerance thresholds [54] for the fault tolerant computation performed by the server. Thus, as long as the preparation on the side of the client can be performed well enough, the correctness of the procedure will ensue. The issue, which cannot be resolved using existing techniques, arises from the other requirement we place on this protocol – blindness.

In many realistic settings, the apparatus the client uses inevitably leaks additional quantum information to the server. A simple example of this case was the setting where the client uses a realistic single photon source which sometimes, albeit rarely, emits two photons instead of one as we discussed previously. Then perfect blindness can no longer be guaranteed, so we need a broader framework to quantify the disturbance done to privacy. This framework we call approximate blindness and is described in this chapter. Within this framework, we also develop a protocol which allows a client to perform UBQC with *arbitrary* levels of security, even if she is restricted to a fixed realistic source of quantum states. We begin by revisiting notions of perfect and approximate secrecy developed for other cryptographic protocols which we build upon.

2.2 Ignorance in two-party schemes: *what it means to have an ignorant server*

We often hear that a particular cryptographic protocol guarantees "privacy", that the eavesdropper "cannot learn anything about our secret", or more generally that a corrupted party cannot "cheat" in any important way. Perhaps the greatest advances in cryptography in the broad sense came about in late 1940's when Claude Shannon introduced a mathematical theory of communication [55], a framework which finally allowed that the vague, but intuitive concepts like secrecy and so on be formally and precisely defined. This is crucial as only then, in a precise framework (of mathematics) can cryptographic guarantees be *mathematically proven*. We begin this chapter by revising some of the definitions and frameworks of what one should consider "ignorance" developed for various classical quantum cryptographic schemes. In these approaches, sensible definitions of "approximate" security, ignorance, or whatever property we may wish for have been defined. Thus, our approach in defining approximate notions of blindness will be to express UBQC in a sensible way in existing frameworks [56, 57, 58] whereby we can inherit the existent notions of "approximate" secrecy and apply it to blindness. What we get is a notion of approximate blindness compatible with similar notions in other cryptographic schemes, which may particularly be of relevance if UBQC is ever to be used as a subroutine in a larger information processing structure [1]. We begin by revising the relevant notions of secrecy.

2.2.1 Perfect secrecy

In (quantum) cryptography, one often considers scenarios in which one party (*e.g* the client in UBQC) has some information which should remain unknown to any other party (for instance, the server in UBQC) even if the latter considers the entire (possibly quantum) system at his disposal. In classical cryptography, statements about the information a party has about a secret in this sense are stated in terms of probability distributions. For instance, let Ξ be a (finite) set of possible secrets, out of which the first party makes a selection. Then, the second party's (the server's) knowledge about the clients choice is represented in terms of a random variable X, taking values in Ξ, with corresponding probabilities p_x, for $x \in \Xi$. The probability $p_x = P(X = x)$ is then understood as the probability the client selected x, *from the perspective of the server.* For illustration purposes let us consider two extremal scenarios. If the server is totally ignorant about the client's choices, then the corresponding probability distribution is uniform - as far as the server is concerned, the client might have picked any secret equally likely. Oppositely, a statement saying that the server has full knowledge that the client picked some particular secret x is represented by the probability distribution $p_x = 1$ and $p_{x'} = 0$, for all $x \neq x'$. In cryptographic protocols, the concealing party often reveals some information, and then the question becomes how this leak is quantified and characterised. Very often the leak is quantified by employing information-theoretical measures like mutual information. Measures are

[1] Ideally, what one would desire is a universally composable definition of blindness. It is an open question if a sensible composable definition can be given which UBQC satisfies. Our approach to approximate blindness will at least guarantee that whatever properties ideal UBQC may have, the approximate variant will inherit, up to a quantifiable probability of failure. More about UBQC and universal composability will be discussed in the next chapters.

again stated in the language of random variables, or probability distributions, but more precisely, *conditional* probability distributions. For instance, let Y be a finite set of messages the client may send to the server, and let the random variable Y take values in Y. Then, if the prior knowledge of the server about the client's secret choice was characterised by the probability distibution $P(X = x)$ the "updated" knowledge about the same secret, once he has received a message $y \in Y$ from the client during the run time of the protocol is represented by the following conditional probability distribution: $P(X = x|Y = y)$. The correlations between the random variables X and Y play a crucial role in determining what sort of harm has been done to the secrecy after additional information is obtained by the server. Again for illustration purposes, we can consider two extremal settings: The random variables X and Y are independent: $P(X, Y) = P(X)P(Y)$. In this case $P(X|Y) = P(X)$ and the additional information did not compromise secrecy. In the other extreme, the random variables may be *perfectly correlated*, for instance let $card(\Xi) = card(\mathsf{Y})$ and let a one-to-one correspondence f between Ξ and Y be known denoted for simplicity as $\mathsf{Y} \ni y_x = f(x)$, for each $x \in \Xi$. Then an example of perfect correlation is represented by the following join distribution $P(X = x, Y = y) = 1/card(\Xi)$ if $y = y_x$ and zero otherwise. In this case after the server received a message $y_{x'}$, this probability distribution is updated as follows: $P(X = x|Y = y_{x'}) = 1$ if $x = x'$ and zero otherwise. In other words, the server knows the client's secret *perfectly*.

In quantum protocols, however, the server (the party we wish to conceal information from) may obtain quantum information–quantum states rather than classical information. The knowledge of the server is then modelled by considering the structure of the joint (bi-partite) state, shared by the server and the client, *as seen by the server*. To illustrate how this works, we will first restate the classical setting considered above in the language of quantum mechanics. First note that the formalism of quantum mechanics where states are represented by positive unit-trace operators directly allows the encoding of classical probability distributions. For example the considered distribution p_x into mixed quantum states:

$$\pi_c = \sum_{x \in \Xi} p_x |x\rangle\langle x|$$

where the states $|x\rangle$ and $|x'\rangle$ are orthogonal for different choices of $x, x' \in \Xi$. Given that the random variable X represents the server's prior state of knowledge about what the client may choose as her secret, this classical state represents the state of the client's system, as viewed from the perspective of the server, after she has chosen her secret. In a more general setting, upon the termination of a protocol between client and server, which ensued the client's choice of a secret, the general state of the system shared by the client and the server may be written as a *classical-quantum* state of the form

$$\pi_{AB} = \sum_{x \in \Xi} p_x |x\rangle\langle x| \otimes \rho_x$$

The state ρ_x represents the state of the server's system, upon the termination of the protocol, *conditional* on the client having chosen x as her secret. Throughout this thesis, we will be denoting the subsystem in the hands of the client with the subscript A, and the server's subsystem

with the subscript B, which are shorthand for Alice and Bob.

Given the state of the shared system π_{AB}, for the server to be ignorant, or oblivious about the client's secret means that the state ρ_x does not depend on x (*i.e.* it is fixed) and that the *a priori* distribution of the random variable X is uniform. In this case, the state π_{AB} is of the form

$$\pi_{AB}^{\text{perfect}} = \left(\frac{\mathbb{1}}{|\Xi|}\right) \otimes \rho^B \tag{2.1}$$

i.e. the system on the client's side is in a totally mixed state and decoupled form the server's system. Above, with $|\Xi|$ we denote the cardinality of the set Ξ. In the classical scenario this corresponds to the setting where the *a-priori* probability distribution about the client's secret choices was uniform, and the probability distribution of the messages the server received from the client was independent from the choice of the secret.

2.2.2 Secrecy with prior knowledge

Recall that we define a secret held by the client to be perfectly secure if the overall system shared by the client and the server, from the perspective of the server, is described with the classical-quantum state

$$\pi_{AB}^{\text{perfect}} = \left(\frac{\mathbb{1}}{|\Xi|}\right) \otimes \rho^B.$$

If the secret corresponds to a variable x chosen from the set Ξ with some *a priori* probability p_x, then the *perfect* state becomes

$$\pi_{AB}^{\text{perfect},\{p_x\}} = \left(\sum_{x\in\Xi} p_x |x\rangle\langle x|\right) \otimes \rho^B \tag{2.2}$$

reflecting the fact that the states he received from the client were independent from the clients choices of the secret, although this secrecy was not perfect to begin with. In other words, this represents the case when the server doesn't learn anything *new* about the client's secret through the information he obtains in a protocol.

2.2.2.1 Approximate secrecy

In practice, however, we often need to work with approximate notions of secrecy, generally defined as follows.

Definition 10. *(informal) We will say that the secret is ϵ-secure (or $\epsilon-$blind as we will call it the context UBQC), with respect to the server's system B if the shared state between the client and the server, π_{AB}, is ϵ-close to the perfectly secure state, $\pi_{AB}^{\text{perfect}}$, defined in 2.1, with respect to the trace distance*

$$\frac{1}{2}\|\pi_{AB}^{\text{perfect}} - \pi_{AB}\|_{\text{tr}} \leq \epsilon\,. \tag{2.3}$$

This notion of secrecy is also convenient since it is well-behaved under composition [58, 59, 57]. Here, we have used the fact that the standard trace norm $\|\cdot\|_{tr}$ induces a metric on the space of density matrices called the trace distance.

Since throughout the remainder of this thesis we will always be working with the trace distance, we will be omitting the subscript tr.

In the framework for characterising security we present, the key object of study is the shared state between the client and the server, as seen by the server. The trace distance is the metric of choice due to it's operational interpretation. It quantifies the probability the server (or any party which is not Alice) can distinguish between the ideal and the realized state. The expression in 2.3 guarantees that the probability of the server (or the environment) successfully guessing, using any type of a procedure allowable by standard quantum mechanics, whether the overall system is the perfect state, or real state is $1/2 - \epsilon/2$, thus, $\epsilon/2$ close to a random guess. Because of this, if Bob chooses a strategy which may result in any type of an information leak, the success probability of this strategy directly depends on ϵ. In other words, the probability ϵ quantifies the probability the protocol's security fails in any possible way. It is worth noting that in this presentation we decouple the security aspect of blindness from correctness (which is a guarantee that the computational output is correct). That is, we will separately give blindness guarantees, and correctness will be guaranteed only in the case of an honest server. This type of decoupling is possible because in this work we deal with a variant of UBQC which does not provide a generic verification of the actions of the server. That is, the only guarantee we care about is that no information is leaked to the server and not whether the computation was performed correctly by the server [2]. The other aspect of delegated computation, called *verifiability*, we will address in Chapter 4.

2.3 Blindness of UBQC revisited

In this section, we will revisit the blindness of UBQC, and attempt to represent the crucial properties of the protocol in such a way that the notions of secrecy and approximate secrecy discussed in previous sections can be applied.

As we explained, the approach adopted in the previous section aimed at explicitly presenting the client's secret and the states the server received throughout the protocol in juxtaposition. A more direct approach is possible. We could also adhere to a modification of the definition of blindness given in the original paper.The alternative approach would be to show that the system the server *has*, at each step of the computation does not depend on the choices of the client, or, for the approximate case, is ϵ close to a state which does not depend on the choices of the client. This approach has been resolved as well and is presented in full form in Chapter 4 Section 4.3. In this

[2]To form an analogy, the one-time pad only guarantees that the eavesdropper learns nothing about the message, while it allows the adversary to alter the message, even in a meaningful way (see the notion of *malleability* in Chapter 6). This security aspect is often called *confidentiality* of channels. If message authentication is also provided, then the channels are called secure. Here, for UBQC, we are predominantly interested in the analogon of confidentiality.

section we proceed to study the systems of the client and the server in conjunction.

In [1], it is shown that the ideal protocol described in the previous section protocol 4 guarantees perfect security with respect to definition 9 which we repeat for the benefit of the reader.

(Blindness) We say a Protocol P *on input* X *is* blind while leaking at most L(X)*, where* $L(X)$ *is any function of the input if:*

1. *The distribution of the classical information obtained by the server in* P *depends only on* $L(X)$.
2. *Given the distribution of classical information described in 1 and* $L(X)$*, the state of the quantum system obtained by the server in* P *is fixed.*

(Blindness theorem) Protocol 4 is blind while leaking at most the dimensions of the brickwork state, i.e. an upper bound on the input size and the depth of the computation.

In this protocol, the client emits a certain number of quantum states, followed by an exchange of classical information. We can then write out the overall quantum state, shared by the client and the server, and proceed with an analysis of this state. Before doing so, it is useful to make some remarks on the UBQC protocol, in order to simplify the form of the shared quantum state.

First, in the original protocol the hidden parameters r_i are chosen by the client during the execution of the protocol. However, they could equally well have been generated in advance. Note also that the only action on the part of the server that affects the behaviour of the client is his reporting of the measurement outcomes, *i.e.* the classical signals s_i. If one investigates how s_i being zero or one affects the client's action (Step 3.5 in **Protocol 4**)[3], one notices that the client's next step, which is sending the measurement angle, only depends on the value of $(s_i + r_i) \mod 2$. Since the random variable r_i is chosen randomly by the client, for any tactic that the server may adopt with respect to the reporting of the signals s_i (*i.e.* for any distribution of the variables s_i), the distribution of the variable $(s_i + r_i) \mod 2$ which actually influences the clients behaviour, is still uniform. More formally, for any sequence of fixed random parameters θ_i and for any sequence of reported signals s_i there exists a sequence of parameters r_i such that the sequence of angles δ_i remains constant. To see this to more detail, we refer the reader to Section 4.3.1, where we meticulously write out the state of the server's system for every configuration of the responses of the server. There, in the final Equation (4.9) we see that the system of the server does not depend on the reporting strategy comprising the possible sequences s_i.

Hence, we can conclude that the server can never influence the client's behaviour. Thus, the sequence of reported signals sent by the server may only jeopardize the correctness of the protocol. As the property of interest in this chapter is blindness, without loss of generality, we may then assume that the signals sent by the server, s_i, are all equal to zero.

In the derivation below, we will assume that the server has no prior information about the computation (parametrized by the choice of the computational angles ϕ_i), *i.e.* that from the server's

[3]The original UBQC protocol does not explicitly model the option of the client to abort. However, what is implicitly assumed is that if the server refuses to send an expected message (the binary signal s_i), then the client will abort.

perspective the *a priori* distribution of the computational angles is uniform, later we will return to this point, and generalize our results for the setting with prior knowledge. The overall quantum state shared by the client and the server in the ideal UBQC protocol running a computation of size S is:

$$\pi_{AB}^{\text{ideal}} = \frac{1}{2^S}\frac{1}{8^S}\sum_{\vec{\phi},\vec{r}}\bigotimes_{i\in[S]}\underbrace{|\phi_i\rangle\langle\phi_i|\otimes|r_i\rangle\langle r_i|}_{\text{Client}}\otimes\underbrace{|+_{\theta_i}\rangle\langle+_{\theta_i}|\otimes|\delta_i\rangle\langle\delta_i|}_{\text{Server}} \tag{2.4}$$

where the classical subsystems $|\phi_i\rangle$ and $|r_i\rangle$ corresponding to the secret, are controlled by the client, and the classical $|\delta_i\rangle$ and quantum $|+_{\theta_i}\rangle$ subsystems are controlled by the server. Here, all the angles are assumed to be in the set $\{k\pi/4\}_{k=0}^{7}$. In the original work on UBQC [1] it was shown that the server cannot access any information about the client's secret (the computational angles ϕ_i or the one-time pad keys r_i) in the case where the overall information is represented by the ideal state (2.4). Here, we will manipulate this state in order to analyse how the server may attempt to access the client's secret and why such an attempt fails.

Because the probability distributions about the hidden parameters r_i and the computational angles ϕ_i are uniform, the entire ideal state factorizes

$$\pi_{AB}^{\text{ideal}} = \frac{1}{2^S}\frac{1}{8^S}\bigotimes_{i\in[S]}\sum_{\phi_i,r_i}|\phi_i\rangle\langle\phi_i|\otimes|r_i\rangle\langle r_i|\otimes|+_{\theta_i}\rangle\langle+_{\theta_i}|\otimes|\delta_i\rangle\langle\delta_i|\,.$$

Here, we have effectively commuted the tensor product and the sum. Then we expand the angles θ_i by reshuffling the definition relation given in Step 3.2 of **Protocol 4**, $\theta_i = \delta_i - \phi_i' + r_i\pi$ (since the angles are computed modulo 2π, the sign in front of $r_i\pi$ is irrelevant, and for simplicity, we will adhere to the convention where it is always set to plus). We then obtain

$$\pi_{AB}^{\text{ideal}} = \frac{1}{2^S}\frac{1}{8^S}\bigotimes_{i\in[S]}\sum_{\phi_i,r_i}|\phi_i\rangle\langle\phi_i|\otimes|r_i\rangle\langle r_i|\otimes|+_{\delta_i-\phi_i'+r_i\pi}\rangle\langle+_{\delta_i-\phi_i'+r_i\pi}|\otimes|\delta_i\rangle\langle\delta_i|\,.$$

Using the classical information δ_i, the server can apply the $Z_{-\delta}$ rotation to his (purely) quantum subsystem, obtaining

$$\pi_{AB}^{\text{ideal}} = \frac{1}{2^S}\frac{1}{8^S}\bigotimes_{i\in[S]}\sum_{\phi_i,r_i}|\phi_i\rangle\langle\phi_i|\otimes|r_i\rangle\langle r_i|\otimes|+_{-\phi_i'+r_i\pi}\rangle\langle+_{-\phi_i'+r_i\pi}|\otimes|\delta_i\rangle\langle\delta_i|\,.$$

Note that, from the perspective of blindness, we can always assume that the server applies any fixed unitary on his part of the system, as a fixed unitary transformation applied locally to the server's part of the system does not alter the correlations between the two subsystems.

Now the subsystem $|\delta_i\rangle$ above no longer carries any relevant information for the server, so we may omit it, and the server can apply the X operator on his quantum system resulting in

$$\pi_{AB}^{\text{ideal}} = \frac{1}{2^S}\frac{1}{8^S}\bigotimes_{i\in[S]}\sum_{\phi_i,r_i}|\phi_i\rangle\langle\phi_i|\otimes|r_i\rangle\langle r_i|\otimes|+_{\phi_i'+r_i\pi}\rangle\langle+_{\phi_i'+r_i\pi}|$$

which we can rewrite as

$$\pi_{AB}^{\text{ideal}} = \frac{1}{2^S}\frac{1}{8^S}\bigotimes_{i\in[S]}\sum_{\phi_i,r_i}|\phi_i\rangle\langle\phi_i|\otimes|r_i\rangle\langle r_i|\otimes Z^{r_i}|+_{\phi_i'}\rangle\langle+_{\phi_i'}|Z^{r_i}\,.$$

In this expression, the server's system is entirely expressed in terms of the adapted computational angles ϕ_i', as described in the last chapter. In the setting where the signals s_i sent by the server are assumed to be zero, the relation between adapted and computational angles are uniquely governed by the parameters r_i about which the server has no information. By ignoring the distinction between the computational and the adapted angles, we are only doing the server a favour and presenting a worse case (in terms of security) scenario for the client (see [1] for a detailed proof of this fact). In the following, we therefore assume that the computational and adapted angles coincide.

Since the server has no information on the hidden parameters r_i, the corresponding subsystem which contains them may then be traced out of the state, and we get

$$\begin{aligned}\pi_{AB/R}^{\text{ideal}} &= \frac{1}{8^S}\bigotimes_{i\in[S]}\sum_{\phi_i}|\phi_i\rangle\langle\phi_i|\otimes\frac{1}{2}\left(|+_{\phi_i'}\rangle\langle+_{\phi_i'}|+|-_{\phi_i'}\rangle\langle-_{\phi_i'}|\right)\\ &= \frac{1}{8^S}\bigotimes_{i\in[S]}\sum_{\phi_i}\left(|\phi_i\rangle\langle\phi_i|\otimes\frac{\mathbb{1}}{2}\right)=\left(\frac{1}{8^S}\bigotimes_{i\in[S]}\sum_{\phi_i}|\phi_i\rangle\langle\phi_i|\right)\otimes\frac{\mathbb{1}}{2^S}\\ &= \frac{\mathbb{1}}{8^S}\otimes\frac{\mathbb{1}}{2^S}\,.\end{aligned}\tag{2.5}$$

Clearly, the server's system is decoupled from the client's, and both are in a totally mixed state. Since we have only been applying unitary operations, restricted to the server's register and independent of the computational angles, this means that the initial state in 2.4, when the secret r parameters have been traced out, attains the same form which defined perfect secrecy. The protocol is therefore perfectly secure, that is, blind. The analysis above only considered the states the server received from the client, and was independent from the actual actions of the server. Hence, we did not assume anything pertaining to actions the server performed, *i.e.* whether he was honest or malicious.

2.3.0.2 Blindness of UBQC under prior knowledge

In the previous subsections, we have addressed the security properties of delegated blind quantum computation protocols under the assumptions that the angles θ_i and the hiding binary digits r_i were chosen uniformly at random, and, more crucially, that the server has no prior knowledge about the computational angles $\overrightarrow{\phi} = (\phi_1,\dots\phi_S)$. The latter condition could potentially be problematic for the following reason. The number of useful algorithms the client might want to perform, or said otherwise, the number of different unitary transformations she wants the server to perform on the state $|+\rangle\dots|+\rangle$, is in general much less than 8^S, which corresponds to the number of possible computational angles choices. Indeed, the vector $\overrightarrow{\phi}$ corresponds to a classical encoding of the unitary transformation U. The encoding map might be *a priori* known

to the server, meaning that the prior probability for a given computation angles vector $\overrightarrow{\phi}$ is not uniform. For this reason, it is essential to establish the blindness of the protocol, even if in the case of the non-uniform distribution of $\overrightarrow{\phi}$. On the other hand, assuming that the random variables θ_i and r_i are initially completely unknown to the server is quite natural since these variables are drawn uniformly at random locally on the client's site.

From now on, we will model the server's prior information about the vector $\overrightarrow{\phi}$ as a probability distribution $p(\overrightarrow{\phi})$. If the delegated computation protocol is blind, then the posterior distribution for the server should be equal to the prior one. This was established for the original UBQC Protocol [1], and in this section we give an alternative proof using the same formalism as previous subsections. We then generalise the notion of ϵ-blindness for UBQC protocols with imperfect client preparation taking into account a non-uniform prior distribution for the computational angles.

Consider now the joint state between the client and the server after all the information has been exchanged in the ideal UBQC protocol, but now assuming prior knowledge about the computational angles given by the distribution $p(\overrightarrow{\phi})$

$$\pi_{AB}^{\text{ideal},\,p(\overrightarrow{\phi})} = \frac{1}{2^S}\sum_{\overrightarrow{\phi},\overrightarrow{r}} p(\overrightarrow{\phi}) \bigotimes_{i\in[S]} \underbrace{|\phi_i\rangle\langle\phi_i| \otimes |r_i\rangle\langle r_i|}_{\text{Client}} \otimes \underbrace{|+_{\theta_i}\rangle\langle+_{\theta_i}| \otimes |\delta_i\rangle\langle\delta_i|}_{\text{Server}}. \tag{2.6}$$

Analogously to the derivation of Section 2.3 for the ideal state without prior knowledge, this state can be rewritten (after tracing out a server subsystem which contains no correlation with the client's variables) as

$$\pi_{AB}^{\text{ideal},\,p(\overrightarrow{\phi})} = \frac{1}{2^S}\sum_{\overrightarrow{\phi},\overrightarrow{r}} p(\overrightarrow{\phi}) \bigotimes_{i\in[S]} |\phi_i\rangle\langle\phi_i| \otimes |r_i\rangle\langle r_i| \otimes Z^{r_i}|+_{\phi_i'}\rangle\langle+_{\phi_i'}|Z^{r_i}$$

where

$$|+_{\phi_i'}\rangle = X^{\left(\oplus_{j\in D_i^X} s_j\oplus r_j\right)} Z^{\left(\oplus_{j\in D_i^Z} s_j\oplus r_j\right)} |+_{\phi_i}\rangle\,.$$

Recall the definitions of the X and Z dependency sets $D_i^{X/Z}$ for the i^{th} state, the phase flip (Z) and the bit flip (X) are conditioned on the hiding parameters r_j with j strictly less than i. This means that the state $\sum_{r_i=0}^{1} Z^{r_i}|+_{\phi_i'}\rangle\langle+_{\phi_i'}|Z^{r_i}$ is the totally mixed state.

Now, since the r_i parameters are unknown to the server, if we trace out the subsystem, denoted R, which contains them, we obtain the state

$$\pi_{AB/R}^{\text{ideal},p(\overrightarrow{\phi})} = \sum_{\overrightarrow{\phi}} p(\overrightarrow{\phi}) \bigotimes_{i\in[S]} |\phi_i\rangle\langle\phi_i| \otimes \frac{\mathbb{1}}{2} = \left(\sum_{\overrightarrow{\phi}} p(\overrightarrow{\phi}) \bigotimes_{i\in[S]} |\phi_i\rangle\langle\phi_i|\right) \otimes \left(\frac{\mathbb{1}}{2}\right)^{\otimes S}.$$

This state corresponds to the perfect state with prior knowledge, equation (2.2), hence the server could learn nothing about the client secret during the execution of the protocol.

2.3.1 Approximate blindness in UBQC

The blindness of the ideal UBQC protocol guarantees that if the overall information shared between the client and the server is represented by the state of the form

$$\pi_{AB}^{\text{ideal}} = \frac{1}{2^S}\frac{1}{8^S}\sum_{\vec{\phi},\vec{r}}\bigotimes_{i\in[S]} \underbrace{|\phi_i\rangle\langle\phi_i| \otimes |r_i\rangle\langle r_i|}_{\text{Client}} \otimes \underbrace{|+_{\theta_i}\rangle\langle+_{\theta_i}| \otimes |\delta_i\rangle\langle\delta_i|}_{\text{Server}}$$

then the server cannot access any information about the client's secrets, which include the computational angles and the encoded input/output state, regardless of the server's actions. Recall that the classical output of the computation, as seen by the server is one-time padded by the hidden parameters r_i, which defined the measurement angles of the final layer of the computation.

As in the proof of the security for the UBQC (see [1]), no assumptions were made on the strategy of the malicious server, since all we considered were the states the server receives, rather than the actual state of his system. The approach where the actions of the server are explicitly modelled we give in Section 4.3.1, for completeness. Thus, there exists no physical transformation, represented by a general completely-positive trace-preserving (CPTP) map $\mathcal{E}$, that can be applied to the server's state (and his private qubits) which would help him learn the client's secrets. We can then define an unconditionally blind state to be, for a given CPTP map $\mathcal{E}$, a state of the form

$$\pi_{AB}^{\mathcal{E}} = \mathbb{1}_{\text{Client}} \otimes \mathcal{E}_{\text{Server}} \left(\frac{1}{2^S}\frac{1}{8^S}\sum_{\vec{\phi},\vec{r}}\bigotimes_{i\in[S]} \underbrace{|\phi_i\rangle\langle\phi_i| \otimes |r_i\rangle\langle r_i|}_{\text{Client}} \otimes \underbrace{(|+_{\theta_i}\rangle\langle+_{\theta_i}| \otimes |\delta_i\rangle\langle\delta_i|)}_{\text{Server}} \right) \quad (2.7)$$

and a family $\mathcal{F}$ of unconditionally blind states as follows

$$\mathcal{F} = \{\pi_{AB}^{\mathcal{E}} \,|\, \mathcal{E} \text{ is a CPTP map}\}.$$

The map $\mathcal{E}$ above signifies any possible deviation from the protocol which can be done by the malicious server. As no deviation will help the server to learn the secret information the client is hiding, hence for any such map the client's secret is unconditionally secure. In equation (2.7), the identity map is applied to the client's subsystem while a general map $\mathcal{E}$ is applied on the server's subsystem. If the overall system shared between the client and the server in a UBQC protocol is given by a state in the family $\mathcal{F}$, then the protocol is blind. If the map $\mathcal{E}$ is the identity, then the shared information between the client and the server corresponds to that of the ideal UBQC protocol, and an honest server can run the UBQC protocol correctly. If this is not the case, the correctness of the executed protocol cannot be guaranteed.

The analyses above characterised the properties of the overall system shared by the client and the server throughout a delegated quantum computation protocol guaranteeing perfect blindness. Next, we focus on protocols that deviate from this ideal setting. This will bring us to a notion of approximate security that is ϵ-blindness. Motivated by the technical difficulties of generating perfect single qubit states for a realistic client, we now consider scenarios in which the client fails to send the exact states of the form $|+_\theta\rangle$ in the preparation phase of **Protocol 4**, but rather some

more general quantum states ρ^{θ} which are in general parametrized by the angle θ. The overall state of the system shared by the client and the server is then given by the following non-ideal state

$$\pi_{AB}^{\{\rho^{\theta_i}\}} = \frac{1}{2^S}\frac{1}{8^S}\sum_{\vec{\phi},\vec{r}}\bigotimes_{i\in[S]}\underbrace{|\phi_i\rangle\langle\phi_i|\otimes|r_i\rangle\langle r_i|}_{\text{Client}}\otimes\underbrace{\rho^{\theta_i}\otimes|\delta_i\rangle\langle\delta_i|}_{\text{Server}}, \tag{2.8}$$

which is characterised by the type of states ρ^{θ_i} emitted by the client. Note that, if the emitted states ρ^{θ_i} are such that there exists a deterministic physical transformation $\mathcal{T}$ performing $\mathcal{T}(\rho^{\theta_i}) = |+_{\theta_i}\rangle\langle+_{\theta_i}|$, the protocol can be made correct by simply demanding that the server applies this transformation on the received states after the client's preparation phase. A UBQC protocol with imperfect client preparation where the emitted states satisfy this property is called *correct*. What remains to be addressed is the blindness of such non-ideal protocols. For this purpose we give the following definition.

Definition 11. *A UBQC protocol with imperfect client preparation, in which in the preparation phase the client sends states of the form ρ^{θ_i} instead of the perfect states $|+_{\theta_i}\rangle$ is called an ϵ-blind UBQC protocol with imperfect states (where $\epsilon \geq 0$ is the security parameter), if the trace distance between the overall joint state given with the expression (2.8) and the family $\mathcal{F}$ of unconditionally blind joint states is less than or equal to ϵ*

$$\min_{\pi_{AB}^{\mathcal{E}}\in\mathcal{F}}\frac{1}{2}\|\pi_{AB}^{\rho^{\theta_i}} - \pi_{AB}^{\mathcal{E}}\| \leq \epsilon\,. \tag{2.9}$$

The criterion above (equation 2.9) can equivalently be written as

$$\min_{\mathcal{E}}\frac{1}{2}\|\pi_{AB}^{\rho^{\theta_i}} - \pi_{AB}^{\mathcal{E}}\| \leq \epsilon$$

where $\mathcal{E}$ ranges over all CPTP maps.

Clearly, if $\epsilon = 0$, such a protocol is equivalent to the ideal UBQC protocol, and is therefore perfectly blind.

This notion of approximate security makes a crucial use of the trace distance between the state obtained while running the actual protocol and an ideal state. This approach has been already used in the context of quantum cryptography and particularly for quantum key distribution, where the ϵ-security of a protocol is defined analogously. The importance of the trace distance comes from the fact that it is closely linked to the maximal probability of distinguishing the actual protocol form the ideal one. If this probability is arbitrary small, then the actual protocol is arbitrary secure. The notion of approximate ϵ-blindness for a UBQC protocol is defined in a natural way in the presence of prior knowledge as well. The ideal blind state, with *a priori* probability distribution $p(\vec{\phi})$ and relative to the CPTP map $\mathcal{E}$, is then defined as

$$\pi_{AB}^{\mathcal{E},p(\vec{\phi})} = \mathbb{1}_{\text{Client}}\otimes\mathcal{E}_{\text{Server}}\left(\frac{1}{2^S}\sum_{\vec{\phi},\vec{r}}p(\vec{\phi})\bigotimes_{i\in[S]}\underbrace{|\phi_i\rangle\langle\phi_i|\otimes|r_i\rangle\langle r_i|}_{\text{Client}}\otimes\underbrace{(|+_{\theta_i}\rangle\langle+_{\theta_i}|\otimes|\delta_i\rangle\langle\delta_i|)}_{\text{Server}}\right).$$

Correspondingly, we define the family $\mathcal{F}^{p(\overrightarrow{\phi})}$ of ideal blind states with respect to prior knowledge $p(\overrightarrow{\phi})$ as follows

$$\mathcal{F}^{p(\overrightarrow{\phi})} = \left\{ \pi_{AB}^{\mathcal{E},p(\overrightarrow{\phi})} \mid \mathcal{E} \text{ is a CPTP map} \right\} .$$

Any protocol characterised by a state in $\mathcal{F}^{p(\overrightarrow{\phi})}$ is as secure as the ideal UBQC protocol with prior knowledge $p(\overrightarrow{\phi})$. Again, like in the setting with no prior knowledge, the correctness of the protocol can be guaranteed only if the map $\mathcal{E}$ is the identity. In general, the state describing a given protocol does not correspond to an ideal state, but is given by the following expression (similar to the state in equation (2.8))

$$\pi_{AB}^{\rho^{\{\theta_i\}_i},\, p(\overrightarrow{\phi})} = \frac{1}{2^S} \sum_{\overrightarrow{\phi},\overrightarrow{r}} p(\overrightarrow{\phi}) \bigotimes_{i\in[S]} \underbrace{|\phi_i\rangle\langle\phi_i| \otimes |r_i\rangle\langle r_i|}_{\text{Client}} \otimes \underbrace{\rho^{\theta_i} \otimes |\delta_i\rangle\langle\delta_i| \otimes \rho^{\theta}}_{\text{Server}} . \tag{2.10}$$

Definition 12. *A UBQC protocol with imperfect client preparation and prior knowledge $p(\overrightarrow{\phi})$, in which the client sends states of the form ρ^{θ_i} instead of the perfect states $|+_{\theta_i}\rangle$ in the preparation phase is called an* ϵ-blind UBQC protocol with imperfect states and prior knowledge $p(\overrightarrow{\phi})$*, if the trace distance between the overall joint state given with the expression (2.10) and the family $\mathcal{F}^{p(\overrightarrow{\phi})}$ of unconditionally blind joint states with prior knowledge is less than ϵ*

$$\min_{\pi_{AB}^{\mathcal{E},p(\overrightarrow{\phi})} \in \mathcal{F}^{p(\overrightarrow{\phi})}} \frac{1}{2} \| \pi_{AB}^{\rho^{\{\theta_i\}_i},\, p(\overrightarrow{\phi})} - \pi_{AB}^{\mathcal{E},p(\overrightarrow{\phi})} \| \leq \epsilon .$$

The above criterion is equivalent to

$$\min_{\mathcal{E}} \frac{1}{2} \| \pi_{AB}^{\rho^{\{\theta_i\}_i},\, p(\overrightarrow{\phi})} - \pi_{AB}^{\mathcal{E},p(\overrightarrow{\phi})} \| \leq \epsilon$$

where the map $\mathcal{E}$ ranges over all CPTP maps.

In the definitions above we have restricted ourselves to a very particular type of states Alice's emiter produces. In the most general setting, Alice may send out states which are mutually entangled, and where there exists classical correlation between the angles θ, which could be caused by an imperfect random angle generator. In this case the notion of approximate blindness would be defined in the same way, operationally speaking, in terms of probability of distinguishing between the states of an ideal system and the system we manage to realize. However, if we focus on preparation, which assumes imperfections as modelled in this section, then we can elegantly characterise the blindness of the entire run of UBQC computation based upon the quality of preparation of individual $|+_{\theta_i}\rangle$ states. In the remainder of this chapter we assume the emitting device has no memory, and that the phase θ choice is driven by a good enough random number generator.

Thus, our goal, for the rest of this section, is to see how an approximate preparation of the states $|+_{\theta_i}\rangle$ affects the blindness of the overall protocol. This is just a theoretical analysis, and an actual

preparation protocol will be given later in this thesis. We wish to bound the distance

$$\min_{\mathcal{E}} \frac{1}{2} \| \pi_{AB}^{\rho^{\{\theta_i\}_i},\, p(\overrightarrow{\phi})} - \pi_{AB}^{\mathcal{E}, p(\overrightarrow{\phi})} \|$$

where the minimization is over all possible CPTP maps $\mathcal{E}$ acting on the system in the possession of the server. We can restrict ourselves to the maps $\mathcal{E}$ which act individually and identically on the subsystems containing the qubits $|+_{\theta_i}\rangle\langle+_{\theta_i}|$

$$\min_{\mathcal{E}} \frac{1}{2} \left\| \pi_{AB}^{\rho^{\{\theta_i\}_i},\, p(\overrightarrow{\phi})} - \pi_{AB}^{\mathcal{E}, p(\overrightarrow{\phi})} \right\| \leq \min_{\mathcal{E}, \text{i.i.d.}} \frac{1}{2} \left\| \pi_{AB}^{\rho^{\theta_i},\, p(\overrightarrow{\phi})} - \pi_{AB}^{\mathcal{E}, p(\overrightarrow{\phi})} \right\| .$$

This clearly holds as the minimization of the right-hand side of the expression above is just the minimization restricted to the subset of the minimization space of the left-hand side of the expression. Omitting the $\frac{1}{2}$ pre-factor, one has

$$\left\| \pi_{AB}^{\rho^{\{\theta_i\}}, p(\overrightarrow{\phi})} - \pi_{AB}^{\mathcal{E}, p(\overrightarrow{\phi})} \right\| =$$

$$\frac{1}{2^S} \left\| \sum_{\overrightarrow{\phi}, \overrightarrow{r}} p(\overrightarrow{\phi}) \left(\bigotimes_{i\in[S]} |\phi_i\rangle\langle\phi_i| |r_i\rangle\langle r_i| \rho^{\theta_i} |\delta_i\rangle\langle\delta_i| - \bigotimes_{i\in[S]} |\phi_i\rangle\langle\phi_i| |r_i\rangle\langle r_i| \mathcal{E}(|+_{\theta_i}\rangle\langle+_{\theta_i}|) |\delta_i\rangle\langle\delta_i| \right) \right\|$$

$$\leq \frac{1}{2^S} \sum_{\overrightarrow{\phi}, \overrightarrow{r}} p(\overrightarrow{\phi}) \left\| \bigotimes_{i\in[S]} |\phi_i\rangle\langle\phi_i| |r_i\rangle\langle r_i| \rho^{\theta_i} |\delta_i\rangle\langle\delta_i| - \bigotimes_{i\in[S]} |\phi_i\rangle\langle\phi_i| |r_i\rangle\langle r_i| \mathcal{E}(|+_{\theta_i}\rangle\langle+_{\theta_i}|) |\delta_i\rangle\langle\delta_i| \right\|$$

$$\leq \frac{1}{2^S} \sum_{\overrightarrow{\phi}, \overrightarrow{r}} p(\overrightarrow{\phi}) \left\| \bigotimes_{i\in[S]} \rho^{\theta_i} - \bigotimes_{i\in[S]} \mathcal{E}(|+_{\theta_i}\rangle\langle+_{\theta_i}|) \right\|$$

$$\leq \frac{1}{2^S} \sum_{\overrightarrow{\phi}, \overrightarrow{r}} p(\overrightarrow{\phi}) \sum_{i\in[S]} \left\| \rho^{\theta_i} - \mathcal{E}(|+_{\theta_i}\rangle\langle+_{\theta_i}|) \right\| . \tag{2.11}$$

Although the variables θ_i are drawn uniformly at random initially, their distribution, given δ_i and some prior knowledge about the angles ϕ_i is not uniform. While the expression (2.11) could be refined further, this general derivation becomes rather cumbersome, and is omitted here. However, if there exists a real value ϵ_{prep} such that

$$\frac{1}{2} \left\| \rho^{\theta} - \mathcal{E}(|+_{\theta}\rangle\langle+_{\theta}|) \right\| \leq \epsilon_{\text{prep}}$$

for all possible angles θ (as it will be the case with the protocol we will present later), we then obtain

$$\begin{aligned} \frac{1}{2} \| \pi_{AB}^{\rho^{\{\theta_i\}_i}, p(\overrightarrow{\phi})} - \pi_{AB}^{\mathcal{E}, p(\overrightarrow{\phi})} \| &\leq \frac{1}{2^S} \sum_{\overrightarrow{\phi}, \overrightarrow{r}} p(\overrightarrow{\phi}) \textstyle\sum_{i\in[S]} \frac{1}{2} \left\| \rho^{\theta_i} - \mathcal{E}(|+_{\theta_i}\rangle\langle+_{\theta_i}|) \right\| \\ &\leq \frac{1}{2^S} \sum_{\overrightarrow{\phi}, \overrightarrow{r}} p(\overrightarrow{\phi}) S \epsilon_{\text{prep}} = S \epsilon_{\text{prep}}. \end{aligned}$$

In such a scenario, the obtained UBQC protocol is ϵ-blind if

$$S\epsilon_{\text{prep}} \leq \epsilon. \tag{2.12}$$

This result implies what one could reasonably hope for: the ability to prepare "high quality qubits" translates to the ability to run UBQC with a high level of blindness.

2.3.1.1 Delegating qubit preparation using coherent light sources

The notion of approximate blindness we introduced here allows us to quantify the blindness levels guaranteed to the client, for a given instrument she uses in the preparation stage. This, however, is not completely satisfying from the client's perspective. Indeed, the client can only achieve a given value of ϵ_{prep} in practice, meaning that for a fixed security parameter ϵ, she cannot perform a computation with more that $\epsilon/\epsilon_{\text{prep}}$ steps. In order to allow for computation of arbitrary size, it is necessary to prepare arbitrary good qubits and the solution is to delegate this task to the server, who is assumed to be much more powerful than the client.

In the next section we will proceed by presenting such a *Remote Blind qubit State Preparation* (RBSP) protocol where the client only needs to prepare weak coherent pulses with a given polarization. The requirements for the client are therefore minimal. In particular, they are the same as in most practical implementations of discrete-variable QKD. The difficulty here is transferred to the server who has to perform a quantum non-demolition measurement to obtain the desired qubit. As we will show, using the RBSP protocol S times, the client can reach a joint state $\pi_{AB}^{\{\rho^{\theta_i}\}}$ which is ε-close to the family $\mathcal{F}$ of perfectly blind states.

Next we present a very brief introduction to some useful concepts from quantum optics we will be using. More details can be found in [11], for instance. The presentation of the notion in quantum optics here is tailored specifically for the purposes of the following section.

The quantum state of a single light mode can be represented as a normed vector in an infinite dimensional Hilbert space called the Fock space. An orthonormal basis of this space, the Fock basis, is given by the number states $\{|i\rangle\}_{i=0}^{\infty}$, the label designating the number of photons occupying that mode. Coherent states are single-mode states characterised by a complex number γ and can be expressed in the Fock basis as follows

$$|\gamma\rangle = e^{-\frac{|\gamma|^2}{2}} \sum_{i=0}^{\infty} \frac{\gamma^i}{\sqrt{i!}} |i\rangle.$$

For a coherent state $|\gamma\rangle$, characterised by the complex number $\gamma = \alpha e^{i\theta}$, the norm $\alpha \geq 0$ is called the real amplitude, and the argument $\theta \in [0, 2\pi]$ the complex phase of the coherent state. The complex phase is only defined with respect to a fixed frame of reference (often called the local oscillator). Coherent states describe the light emitted by high-quality lasers. The real amplitude α is directly correlated to the average photon number $\langle n \rangle$ occupying the mode

$$\alpha^2 = \langle n \rangle.$$

For our purpose in the context of UBQC, the phase of the coherent state will not matter, and we will in general consider randomised version of the coherent state

$$\rho(\alpha) = \frac{1}{2\pi}\int_0^{2\pi} |\alpha e^{i\theta}\rangle\langle\alpha e^{i\theta}| d\theta = e^{-\alpha^2}\sum_{i=0}^{\infty}\frac{\alpha^{2i}}{i!}|i\rangle\langle i|.$$

In this expression, the state $|i\rangle$ corresponds to i copies of a single photon. In particular, if the polarisation of the coherent state is $|+_\theta\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta}|1\rangle)$, then one has $|i\rangle := |+_\theta\rangle^{\otimes i}$. In the following, we will denote $|i\rangle\langle i|_\theta$ the state corresponding i photons with polarisation θ. If $i = 0$, the state is the vacuum state.

2.4 Remote blind qubit state preparation

The results of the previous section have shown that the ability to prepare good approximations of the states $|+_{\theta_i}\rangle$ provably translates into the ability to perform approximately-blind universal quantum computing.

As mentioned, this is not completely satisfying from the client's perspective as in order to allow for computation of arbitrary size it is necessary to prepare "arbitrary good" qubits, for a fixed level of security. We resolve this issue by delegating the difficult preparation task to the more powerful server, which a client can run with only reasonable powers.

We proceed by presenting a *Remote Blind qubit State Preparation* (RBSP) protocol where the client only needs to prepare weak coherent pulses with a given polarization. The requirements for the client are therefore minimal. In particular, they are the same as in most practical implementations of discrete-variable QKD. The difficulty here is transferred to the server who has to perform a quantum non-demolition measurement to obtain the desired qubit. As we will show, using the RBSP protocol S times (which is the number of qubits needed for a particular computation), the client can reach a joint state $\pi_{AB}^{\{\rho^{\theta_i}\}}$, used in the approximate blindness framework, which is ε-close to the family $\mathcal{F}$ of perfectly blind states. This allows for ε-blind UBQC for any value ε of client's choosing.

The RBSP protocol is thus designed to serve as a substitute for the process of sending one individual perfect random qubit which allows for imperfect devices and channel.

Let us elaborate on the relevant characteristics of the process in which a client sends a perfect random qubit to the server that comprises the client's preparation phase. The client generates a state $|+_\theta\rangle\langle+_\theta|$ where the angle θ is chosen uniformly at random from the set of the eight allowable angles. This state is emitted to the server, and is never lost or perturbed. At this point, the only system held by the server which is correlated to the value θ, kept by the client is the qubit state $|+_\theta\rangle\langle+_\theta|$. A malevolent server may at that point perform any physical transformation on the qubit he received. This however, may only jeopardize the correctness of the ensuing UBQC protocol, but not the blindness.

Thus, the end result of the ideal process we wish to emulate is characterised by the following

properties: **(A)** the state in the server's possession is $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$ for a CPTP map $\mathcal{E}$, independent of θ known to the client alone – guaranteeing perfect blindness; **(B)** the protocol is never aborted in the honest server scenario – guaranteeing the robustness of the encompassing UBQC; **(C)** in the honest server scenario, the map $\mathcal{E}$ is the identity (when no imperfections are present) – guaranteeing the correctness of the UBQC protocol.

The RBSP protocol using weak coherent pulses, which we now present, approaches the properties above asymptotically: using it, the client achieves ϵ-blindness, as described in the previous section, and ϵ-robustness guaranteeing that the honest abort probability is less than ϵ. As we will show, the client can achieve an arbitrarily small value of ϵ efficiently by tuning a parameter of the protocol as we explain presently. Despite the imperfect preparation stage, we also show that the correctness of the protocol holds in the honest scenario, whenever the client does not abort (which occurs only with ϵ probability).

To run the RBSP protocol, the client sends a sequence of N weak coherent pulses (small amplitude, phase-randomized coherent states) with random polarization σ in the set $\{0, \pi/4, \ldots, 7\pi/4\}$ to the server. If the transmittance of the channel from the client to the server is supposed to be at least T, then the mean photon number of the source is set to $\mu = T$. This value of $\mu = T$ is optimal for our security analysis, however, other values are in principle admissible as well.

The introduced phase randomization simplifies the security analysis and causes the state emitted from the source to be:

$$\rho^\sigma = \sum_{k=0}^{\infty} p_k |k\rangle\langle k|_\sigma$$

where $|k\rangle_\sigma := |+_\sigma\rangle^{\otimes k}$ corresponds to k photons, occurring with probability $p_k = e^{-T}T^k/k!$, with polarization σ. Each pulse is then a probabilistic mixture of Fock states. The Poissonian distribution obtained here is not crucial for the RBSP protocol. For instance, it would work equally well (with re-adjustment of parameters) with any source realizing a mixture of polarization encoded photon number states, such as polarized thermal states, provided that the probability of getting a single photon is not too small.

The server then performs non-demolition photon number measurements on the pulses he receives, declaring the number outcomes to the client. This additional requirement on the quantum server of non-demolition photon counting, while a challenging task, has already witnessed first experimental demonstrations [60]. At this point, the client checks the number of reported vacuum states - if this number is greater than $N(e^{-T^2} + T^2/6)$, she aborts the protocol. A higher value would be indicative to either a lossier than believed channel, or more importantly, that the server lied in an attempt to cheat.

Before we proceed, let us consider the states the server has received from the client at this point. The server received N phase-randomized coherent pulses, with random polarizations known to the client alone, and through photon counting, these are translated to N collections of qubits, each collection with a secret polarization σ. With a probability exponentially high in N, at least

one of these collections comprises exactly one qubit in the state known to the client alone. This "single-copy" would be the ideal qubit required for UBQC. However, the client cannot control *which* of the N coherent pulses will be measured as a single photon, or trust the server that he shall report honestly on this event. The solution to this problem comes by requiring the server to perform a subroutine using all the qubits he has received which serves as a sort of "blindness amplification". The required subroutine has the property that its resulting qubit state depends equally and completely on each qubit it receives on input – then, if *any* of the qubits used in this subroutine is such that the state of this qubit is unknown to the server, so is the final output. The subroutine we propose we call the interlaced 1D cluster computation subroutine (I1DC).

Thus, if the client is certain that at least one of the qubits the server uses in I1DC is such a "single-copy" qubit, then she can rest assured that the final output is a perfectly "blind" qubit. The threshold $N(e^{-T^2}+T^2/6)$ on the allowable number of reported vacuums serves to guarantee just this.

Continuing with the RBSP protocol description, provided the protocol was not aborted after the server has announced the received number of vacuums, the server performs the I1DC subroutine using the photons obtained by the number measurement of the received coherent pulses.

In this subroutine, the server reports the received photon numbers for each pulse and couples the first and the second qubit (*i.e.* photons) with the interaction $\wedge Z.(H \otimes \mathbb{1})$, and the first qubit of the pair is then measured in the Pauli X basis and the measurement outcome is sent to the client. The remaining qubit is then coupled to the third qubit in the input set and measured in the same basis. This process is repeated until only one qubit remains unmeasured, in some state $|+_\theta\rangle$.

Using her knowledge about the polarizations of each of the pulses initially sent, the declared photon number measurement outcomes, and the reported binary string of outcomes in the I1DC subroutine, the client can compute the angle θ as shown in the algorithm 5. The pair θ (held by the client) and $|+_\theta\rangle$ (held by the server), is the required outcome of the RBSP protocol.

The intuition behind this protocol is the following. The I1DC subroutine is such that if the server is totally ignorant about the polarization of at least one photon in the 1D cluster, then he is also totally ignorant about the final angle θ. Intuitively, this can be seen from the form of the output state in Equation 2.13; the angle θ the client will store as her secret angle is the sum of all the polarization angles the server receives. If one of the angles in this sum comes from a pulse which was measured to contain only one photon, then, whatever the server does, he cannot learn more about θ then he could by first following the protocol honestly, and then by inspecting the resulting state (which can also be seen as a single polarized photon). But, if the server follows the protocol and obtains the resulting single qubit state, then the protocol achieved its goal.

In order to exploit this property, the client should make sure that the server will at least once measure a single photon and put it in the cluster. The cheating strategy for the server consists in claiming he received 0 photon when he received 1 and claiming he received 1 when he in fact measured several (in which case he can learn something "extra" about their polarization). In order to avoid this attack, the client simply verifies that the reported statistics of the server are compatible with the assumed transmittance of the channel. This we elaborate further later,

where we also formally prove the security requirements of RBSP for all malicious activities of the server.

The Interlaced 1D cluster computation protocol, used as a subroutine in RBSP, is described in Protocol 5.

Protocol 5 Interlaced 1-D Cluster computation (I1DC)

1. **Input:** A sequence of k states $(|+_{\sigma_l}\rangle)_{l=1}^{k}$ for $\sigma_l \in_R \{\frac{j\pi}{4}\}_{j=0}^{7}$.

2. **Output:** A binary string of measurement outcomes $s = (s_1, \ldots, s_k)$ and the state $|+_\theta\rangle$, where

$$\theta = \sum_{l=1}^{k} (-1)^{t_l} \sigma_l \,, \tag{2.13}$$

   where the binary components t_i are given as follows:

$$t_i = \begin{cases} \sum_{j=i}^{k-1} s_i \mod 2, & \text{for} \quad i < k \\ 0 & \text{for} \quad i = k \end{cases} \tag{2.14}$$

3. **Computation steps:**
   (a) For $i = 1$ to $(k-1)$
      i. Apply the unitary $\wedge Z.(H \otimes \mathbb{1})$ to qubits i and $i+1$.
      ii. Measure qubit i in the Pauli-X basis, obtaining the outcome s_i.
   (b) Return the string $s = (s_1, \ldots, s_k)$ and the remaining non-measured qubit in the state $|+_\theta\rangle$.

The complete RBSP protocol using phase-randomized weak coherent pulses is given in Protocol 6.

The performance of the described protocol asymptotically approaches the properties of the process of sending ideal qubits in the ideal setting we described at the beginning of this section. More precisely, we have that for the described RBSP protocol, property **(A)** holds except with probability p_{fail} and properties **(B)** and **(C)** hold except with probability p_{abort}. These probabilities p_{abort} and p_{fail} can be bounded as functions of the transmittance T and the parameter N as follows:

$$p_{\text{fail}},\, p_{\text{abort}} \leq \exp\left(-\frac{NT^4}{18}\right) . \tag{2.16}$$

While in the intuitive explanations we described thus far we have assumed the honest behaviour of the server, the properties listed above hold for any malevolent strategy. We prove this formally in the Section 2.6, and here give only the outline of the proof.

The proof of these claims comprises roughly four main parts.

First, we may immediately note that we can always assume that the server always performs the requested photon number measurement, as this measurement operator commutes with the density operator of the states he receives from the client. Thus, we will be concerned with the possible reporting strategies of the server, namely his choice of *reported* photon numbers, versus the true

Protocol 6 Remote Blind qubit State Preparation with weak coherent pulses with parameters (N, T)

1. **Client's preparation**
   (a) The client generates N weak coherent pulses with mean photon number $\mu = T$ and a randomized phase and a polarisation σ_l (for $l = 1, \ldots, N$). These states are described by
   $$\rho^{\sigma_l} = e^{-\mu} \sum_{k=0}^{\infty} \frac{\mu^k}{k!} |k\rangle\langle k|_{\sigma_l}. \tag{2.15}$$
   The polarisation angles σ_l are chosen uniformly at random in $\{k\pi/4 : 0 \leq k \leq 7\}$. The client stores the sequence $(\sigma_1, \ldots, \sigma_N)$.
   (b) The client sends the states $\{\rho^{\sigma_l}\}_l$ to the server.

2. **Server's preparation**
   (a) For each state he receives, the server performs a non-demolition measurement of the photon number, obtaining a sequence of N classical values and N post-measurement states. If the measured photon number was greater than zero, the server keeps one photon, discarding the rest.
   (b) The server reports the string $(n_1, \ldots, n_N)$ to the client.

3. **Client-server interaction**
   (a) The client verifies that the reported number of vacuum states is not too large with respect to the tolerated value of the transmittance T of the quantum channel between the client and the server. More precisely, if this number is larger than $N(e^{-T^2} + T^2/6)$, then the client aborts the protocol.
   Otherwise the protocol continues.
   (b) The server discards the systems for which he measured zero photon. Each subsystem with $n_l > 0$ photons measured, parametrized by the polarisation σ, is interpreted as a system of n_l qubits in the state $|+_{\sigma_l}\rangle$. Only one qubit copy per received state is kept, and the total remaining number of qubits is M.
   (c) Using the qubits from the step above and respecting the sending order, the server performs the I1DC computation (see **Protocol 5**), obtaining the sequence $t = (t_1, \ldots, t_M)$ and keeping the resulting state $|+_\theta\rangle$.
   (d) The server reports the string t to the client.
   (e) Using her knowledge about the angles σ_l of the qubits used in the I1DC procedure by the server, and the received outcome string t, the client computes θ with formula (2.13).

measured photon numbers.

Following this we proceed to prove the claim **(C)** above, that is if the client did not abort and the server is honest, then the server has a "single-copy" qubit in the state $|+_\theta\rangle$ where θ is known to the client alone. This is a simple consequence of the correctness of the I1DC subroutine, proven in Section 2.6, and the fact that the polarization angles are chosen uniformly at random.

Next, we focus on two values: 1) the honest setting abort probability p_{abort} and, 2) the probability p_{fail} of the Protocol *not* aborting, for a particular malevolent reporting strategy in which the server always reports zero when he measured zero or one photons, and reports one or more when he received one or more photons. For these probabilities we derive the bounds stated above using statistical techniques. In particular we use the Hoeffding's inequalities which bounds the probability that an empirical mean of a repeated random experiment deviates from the true expectation by more than a particular threshold, as a function of the number of repeats of the experiment (in our case, this will be the pulse number N).

Finally, we prove that for any other reporting strategy, and any other type of deviation on the part of the server, which includes *not* performing the I1DC subroutine, but anything else allowed by quantum mechanics, the state in his system always attains the from $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$ for a CPTP map $\mathcal{E}$ where the map $\mathcal{E}$ is the contraction to the state in the server's possession, independent of the client's angle θ she computed based on the servers reports. This is shown by two technical lemmas we state and prove at the end of Section 2.6.

2.5 Blind quantum computing with weak coherent pulses

We have broken up UBQC into a preparation phase and the computation phase, and now we put the protocol back together.

As we discussed before, the only source of imperfection which can jeopardize blindness in the implementation of the UBQC protocol comes from the fact that the client cannot generate the states $|+_\theta\rangle$ exactly, but only an approximate state ρ^θ. Using the remote blind qubit state preparation protocol with parameters (N, T) the state generated instead of $|+_\theta\rangle$ can be described, in a worst-case scenario from the client's point of view, as

$$\rho^\theta = (1 - p_{\text{fail}})\mathcal{E}^S(|+_\theta\rangle\langle+_\theta|) + p_{\text{fail}}|\theta\rangle\langle\theta|, \tag{2.17}$$

where $p_{\text{fail}} \leq \exp\left(-\frac{NT^4}{18}\right)$ and $|\theta\rangle$ is a classical state giving full information about θ [4]. We note that in the analysis we presented we have assumed no errors in the angle preparation of the client. The effect of such imperfections is addressed later in the discussion Section 2.7.

This state corresponds to the worst case scenario because we assume that if the RBSP protocol

[4]The state space of the sub-register containing the state $\mathcal{E}^S(|+_\theta\rangle\langle+_\theta|)$ or the state $|\theta\rangle\langle\theta|$ above is of dimensionality no less than eight (to be able to store the classical information) and both of the states should be thought of as being encoded using some orthonormal basis of the register. For instance, if the register state space is spanned by $\{|k\rangle\}_{k=0}^8$ then the state $|\theta\rangle$ could be encoded as $|\theta\rangle := |k\rangle$ for $\theta = k\pi/4$ and the qubit state $|+_\theta\rangle$ as $|+_\theta\rangle := 1/\sqrt{2}(|0\rangle + \exp(ik\pi/4)|1\rangle)$

fails, which happens with probability p_{fail} at most, then the server obtains complete information about the angle θ chosen by the client. As we have shown, in any other cases the server has the state $\mathcal{E}^S(|+_\theta\rangle\langle+_\theta|)$ for some CPTP map $\mathcal{E}^S$, independent of the angle θ computed by the client. We have that

$$\frac{1}{2}\left\|\rho^\theta - \mathcal{E}(|+_\theta\rangle\langle+_\theta|)\right\| \leq p_{\text{fail}}$$

for every θ and the CPTP map $\mathcal{E}(\rho) = \mathcal{E}^S(\rho)$. Hence, in order to characterise blindness, one just needs to consider the condition described in equation (10.77), that is $S\epsilon_{\text{states}} \leq \epsilon$ and to note that

$$p_{\text{fail}} \leq S \exp\left(-\frac{NT^4}{18}\right).$$

The latter claim holds as we have taken into account that the remote blind qubit state preparation protocol is used S times during the complete UBQC protocol. The probabilities that blindness or robustness is jeopardized during this whole process can be bounded easily with the union bound as they are simply increased by a factor S. Hence, choosing $N = 18\ln(S/\epsilon)/T^4$ allows the client to obtain ϵ-blindness for the overall UBQC protocol for arbitrary small values of ϵ. The above explanation, given the proven bounds on the probabilities p_{fail} and p_{abort}, proves our main theorem:

Theorem 5. *A UBQC protocol of computation size S, where the client's preparation phase is replaced with S calls to the coherent state Remote Blind qubit State preparation protocol, with a lossy channel connecting the client and the server of transmittance no less than T, is correct, ϵ-robust and ϵ-blind for a chosen $\epsilon > 0$ if the parameter N of each instance of the Remote Blind qubit State preparation protocol called is chosen as follows:*

$$N \geq \frac{18\ln(S/\epsilon)}{T^4}. \tag{2.18}$$

We note that in the Theorem above we do not assume an upper bound on the transmittance, but only a lower bound.

2.6 Details of security proofs

In this section we give the detailed proofs of claims which were too technical to be included in the main part of this chapter.

Lemma 13. *Protocol 5 is correct.*

Proof. We will prove by induction that the state of the output qubit in the interlaced 1-D computation protocol performed on the input of k qubits in the states $\{|+_{\sigma_l}\rangle\}_{l=1}^k$, given the sequence of measurement outcomes $(s_1, \ldots, s_{k-1})$ is the state $|+_\theta\rangle$, where

$$\theta = \sum_{l=1}^{k}(-1)^{t_l}\sigma_l \tag{2.19}$$

where the binary parameters $(t_1, \ldots, t_k)$ are computed as follows:

$$t_i = \begin{cases} \sum_{j=i}^{k-1} s_i \mod 2, & \text{for} \quad i < k, \\ 0 & \text{for} \quad i = k. \end{cases} \tag{2.20}$$

For the basis of the induction we verify that the claim holds for the first non-trivial case, $k = 2$. Consider the state

$$\wedge Z(H \otimes \mathbb{1})|+_{\sigma_1}\rangle \otimes |+_{\sigma_2}\rangle.$$

It is easy to check that the state of the second subsystem, after the measurement of the Pauli-X observable on the first subsystem is the state $|+_{\sigma_2+(-1)^{s_1}\sigma_1}\rangle$, where $s_1 = 0$ corresponds to the measurement outcome associated to the post-measurement state $|+\rangle$ and $s_1 = 1$ to the outcome associated to the post-measurement state $|-\rangle$. Then, according to equation (2.14), $t_1 = s_1$ and $t_2 = 0$, and then equation (2.13) gives

$$\theta = (-1)^{t_1}\sigma_1 + (-1)^{t_2}\sigma_2 = (-1)^{s_1}\sigma_1 + \sigma_2,$$

which is the angle corresponding to the resulting state for the case $k = 2$.

Assume then the step of the induction, *i.e.* that the claim holds for the input size $k = n$, and let us show that it then also holds for $k = n + 1$. Consider the case where the computational steps of the I1DC protocol have been run to the n^{th} step, i.e. to finish off the protocol, the output qubit of the first n steps of the computation needs to be entangled to the $(n+1)^{st}$ qubit using the prescribed interaction and measured in the Pauli X eigenbasis. Let $(s_1, \ldots, s_{n-1})$ be the measurement outcomes of the first $n - 1$ measurements. Then by the step of the induction the state of the output qubit of the first n steps is $|+_{\theta'}\rangle$ where

$$\theta' = \sum_{l=1}^{n} (-1)^{t'_l}\sigma_l$$

and

$$t'_i = \begin{cases} \sum_{j=i}^{n-1} s_i \mod 2, & \text{for} \quad i < n, \\ 0 & \text{for} \quad i = n. \end{cases}$$

If the entangling interaction is then applied on this resulting qubit $|+_{\theta'}\rangle$ and the remaining qubit $|+_{\sigma_{n+1}}\rangle$, and the first qubit is then measured in the Pauli X eigenbasis, by the basis of the induction, the resulting state is $|+_\theta\rangle$ where:

$$\theta = (-1)^{s_n}\theta' + \sigma_{n+1}.$$

This in turn can be expanded as

$$\begin{aligned}
\theta &= (-1)^{s_n}\theta' + \sigma_{n+1} \\
&= (-1)^{s_n}\sum_{l=1}^{n}(-1)^{t'_l}\sigma_l + \sigma_{n+1} \\
&= (-1)^{s_n}\sum_{l=1}^{n}(-1)^{\left(\sum_{j=i}^{n-1} s_i \mod 2\right)}\sigma_l + \sigma_{n+1} \\
&= \sum_{l=1}^{n}(-1)^{\left(\sum_{j=i}^{n} s_i \mod 2\right)}\sigma_l + \sigma_{n+1} \\
&= \sum_{l=1}^{n+1}(-1)^{t_l}\sigma_l
\end{aligned}$$

for

$$t_i = \begin{cases} \sum_{j=i}^{n-1} s_i \mod 2, & \text{for} \quad i < n+1, \\ 0 & \text{for} \quad i = n+1. \end{cases}$$

Hence, the I1DC protocol is correct. □

Security analysis of remote blind qubit state preparation

The Remote Blind qubit State Preparation protocol (RBSP) is described by Protocol 6. In order to characterise the security characteristics of RBSP, we show the following properties which together with the properties of the original UBQC protocol in [1] prove the claims stated in the main Section 2.4:

(A) Upon the completion of the RBSP protocol the state in the server's possession is $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$ for some CPTP map $\mathcal{E}$ (independent of θ) and the client alone knows the angle θ, except with probability p_{fail};

(B) The protocol is never aborted in the honest server scenario, except with probability p_{abort};

(C) In the honest server scenario, the map $\mathcal{E}$ is the identity if the client did not abort and the protocol is correct.

The correctness of the protocol in property **(C)** above means that upon the successful completion of the RBSP protocol, the server has the state $|+_\theta\rangle\langle+_\theta|$ where θ is the angle the client has computed.

We claim that the probabilities p_{fail} and p_{abort} are bounded above in terms of the protocol parameter N, and relative to the transmittance lower bound T as follows:

$$p_{\text{fail}},\, p_{\text{abort}} \leq \exp\left(-\frac{NT^4}{18}\right).$$

Proof. We begin by proving **Claim (C)**, which is a consequence of the correctness of the interlaced 1-D computation protocol, Lemma 13. For **Claim (C)** to hold, first it needs to be shown

that if the protocol was not aborted, and the server is honest, then the server's system is in the state $|+_\theta\rangle$ for some θ known only to the client. In the case where the server is honest, prior to the call to the interlaced 1D cluster computation subroutine, the server's system is in the state

$$\bigotimes_{l=1}^{k} |+_{\sigma_k}\rangle \tag{2.21}$$

where the angles σ_k are known to the client. Then the server will perform the interlaced 1D cluster computation using this system as the input, reporting the bit string $(t_1, \ldots, t_k)$, which is related to the measurement outcomes, as explained in Protocol 5. The client will then compute the angle θ using the formula (2.13). Hence, by Lemma 13 this angle is precisely the angle defining the state $|+_\theta\rangle$ in the server's subsystem. What remains to be seen is that the angle of this resulting state is chosen uniformly at random. Recall that the angle θ is given with $\theta = \sum_{l=1}^{k}(-1)^{t_l}\sigma_l = \sum_{l=1}^{k-1}(-1)^{t_l}\sigma_l + \sigma_k$, and since the angles σ_k are polarisation angles and they are assumed to be chosen uniformly at random, the angle θ is also distributed uniformly at random. This proves that the client alone knows the value of θ.

To prove **Claim (B)**, we will need to bound the abort probability when the server is honest. Finally, for **Claim (A)**, we will need to show that if the protocol is not aborted then the state in the server's possession is $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$ for some CPTP map $\mathcal{E}$, where θ is the angle the client will compute based on the servers feedback, except with probability p_{fail}.

We address these two required properties throughout the rest of this section. In Lemma 14, which we present later, we show that if during the run-time of the protocol the server measures a single photon in one of the states (coherent pulses) sent by the client and declares it as such, then if the client does not abort the protocol, the resulting state with the server is $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$ for a CPTP map $\mathcal{E}$. Here, θ is the angle the client will compute based on the servers feedback. Hence, the probability of this not happening, is the failure probability, p_{fail}.

Here, we note $p_k = e^{-\mu}\frac{\mu^k}{k!}$ the probability of receiving k photons if the channel is perfect (unit transmittance), and $p_k^T = e^{-T\mu}\frac{(T\mu)^k}{k!}$ the probability of receiving k photons if the quantum channel between the client and the server is a lossy channel of transmittance T, where μ denotes the mean photon number of the coherent pulse. In fact, since the events with 2 or more photons are not distinguished by our protocol, we note $p_{\geq 2}$ (resp. $p_{\geq 2}^T$) the probability of obtaining 2 or more photons for a perfect channel (resp. a channel with transmittance T).

In what follows, we derive the bounds for both p_{fail} (blindness) and p_{abort} (robustness). For each state that the server receives, he is supposed to perform a non-demolition measurement of the photon number and to announce this number to the client. Here, we are only interested in three types of events:

- "event 0" when the server measures 0 photon. This event has probability p_0^T in the case of an honest server since the transmission channel is characterised by a transmittance T.
- "event 1", when the server measures exactly 1 photon. The whole point of the protocol is to make sure that at least once, this event occurs and the server has to announce that he

received one photon. If this is the case, by Lemma 14 we are guaranteed the server has the desired state $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$ for a CPTP map $\mathcal{E}$.

- "event 2", when the server measures at least 2 photons. In the case of a malicious server, one has to suppose that the probability of such an event is $p_{\geq 2}$ (instead of $p^T_{\geq 2}$), meaning that we assume that the server has the ability to replace the imperfect quantum channel by a lossless one.

As we will explain below, without loss of generality we may assume that the server always performs the number measurement. Then, if the server is malicious, his only strategy consists in declaring he received 0 photon when he detected 0, declaring he received 0 photon when he detected 1, and declaring either 1 or more when he detected at least 2 photons [5]. Any other strategy will either mean that the server will admit to having measured one photon in which case, by Lemma 14 the protocol will end in a satisfactory state. Alternatively, the server has to report that he measured 1 photon when he in fact measured none. In this case in the setting where the client did not abort, the angle the client computes will be uncorrelated to the state generated by the server. This will compromise the correctness of the computation, however, it will corresponds to a state of the form $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$ for a CPTP map $\mathcal{E}$ where the map $\mathcal{E}$ is the contraction to the state in the server's possession, independent of the client's calculated angle θ. Thus, it does not jeopardize our property of interest – blindness. We will prove this formally in Lemma 15 presented later.

Let us denote with N the total number of states sent by the client, M_0, M_1 and M_2 the number of states for which the server *measured* respectively 0, 1 or at least 2 photons. Also define N_0, N_1 and N_2 to be the respective numbers of states for which the server *reported* having measured 0, 1 or at least 2 photons. Note that the numbers M_0, M_1 and M_2 are well-defined since the server does not gain anything by not measuring the photon number for each state he receives. This is because the measurement operators commute with the state sent by the client, which are diagonal in the Fock basis. We can therefore assume that he performs this non-demolition measurement.

These various quantities are related through the normalization constraint

$$M_0 + M_1 + M_2 = N_0 + N_1 + N_2 = N.$$

For an honest server, one has $N_0 = M_0, N_1 = M_1, N_2 = M_2$. A malicious server will, however, choose a strategy such that $N_0 = M_0 + M_1$. Consider the probability that the protocol aborts when the server is honest. Hoeffding's bound [61] immediately gives an upper bound for p_{abort}, for any $\Delta > 0$ we have

$$\begin{aligned} p_{\text{abort}} &= \Pr\left[\frac{M_0}{N} - p_0^T \geq \Delta\right] \\ &\leq \exp(-2\Delta^2 N). \end{aligned} \tag{2.22}$$

[5] Here, by *strategy* we mean the strategy pertaining to the declared photon numbers. We do not assume anything about what the server may do following his declaration.

The only way the protocol fails is that the malicious server applies the strategy described above, (that is, to pretend he did not receive anything unless he actually received at least two photons) while not being detected. Let us consider a tolerance Δ which will be optimized later. One has:

$$\begin{aligned} p_{\text{fail}} &= \Pr\left[\frac{N_0}{N} - p_0^T \leq \Delta\right] \\ &\leq \Pr\left[\frac{M_0 + M_1}{N} - p_0^T \leq \Delta\right] \\ &\leq \Pr\left[1 - \frac{M_2}{N} - p_0^T \leq \Delta\right] \\ &\leq \Pr\left[\frac{M_2}{N} - p_2 \geq 1 - p_0^T - p_{\geq 2} - \Delta\right] \\ &\leq \exp(-2\tilde{\Delta}^2 N), \end{aligned} \tag{2.23}$$

with

$$\tilde{\Delta} := 1 - p_0^T - p_{\geq 2} - \Delta.$$

In order to get a non-trivial bound for p_{fail}, the parameter $\tilde{\Delta}$ should be positive and bounded away from 0. One has

$$\begin{aligned} \tilde{\Delta} + \Delta &= 1 - e^{-T\mu} - (1 - (1+\mu)e^{-\mu}) \\ &= e^{-\mu}\left(1 + \mu - e^{(1-T)\mu}\right). \end{aligned}$$

If we fix $\mu = T$, we obtain

$$\tilde{\Delta} + \Delta = e^{-T}\left(1 + T - e^{T(1-T)}\right) \geq \frac{T^2}{3}.$$

Hence, choosing $\Delta = \tilde{\Delta} \geq T^2/6$, one gets

$$p_{\text{fail}},\, p_{\text{abort}} \leq \exp\left(-\frac{NT^4}{18}\right).$$

□

While the probabilities p_{fail} and p_{abort} can in principle be made arbitrary small for any (positive) value of the transmittance, one notes that the required number of weak coherent pulses scales like $\log(1/\epsilon)/T^4$ for small T, making the scheme less efficient. In general, this subroutine will be used S times during the complete UBQC protocol. The probabilities that blindness or robustness is jeopardized during this whole process can be bounded easily with the union bound and they are simply increased by a factor S. This means that the correct scaling for the parameter N should be $(\log(S/\epsilon))/T^4$.

In the above, we considered one specific implementation of the remote blind qubit state preparation protocol using weak coherent pulses. This choice was made because weak coherent pulses are arguably the simplest quantum states to prepare in a laboratory. However, the protocol could be easily generalised to any source of light that emits a mixture of Fock states. In particular, the

protocol would work equally well with a thermal source of light. The only characteristics which are required are that the probability of emitting exactly one photon is strictly positive and that the client is able to calibrate her source well enough. In other words, reasonable bounds on the probability of emitting a given number of photons should be available.

Next, we present the lemmas we need to complete the proof above.

Lemma 14. *If the server measured a weak coherent pulse sent by the client to contain one photon, declared it as such honestly to the client, and the client did not abort in the presented remote blind qubit state preparation protocol, then the state in the possession of the server after the termination of the protocol is $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$ for a CPTP map $\mathcal{E}$ where θ is the angle computed by the client.*

Proof. We begin the proof by describing the system of the client and the server after the server has reported the binary string $\{t_i\}_i$ to the client, which he has to do to prevent the client from aborting. The client has the following:

- A sequence of angles $\{\sigma_k\}_{k=1}^M$ which the client has encoded in the polarization of the initially sent coherent pulses, corresponding to those pulses for which the server has announced a non-zero declared photon number. In this sequence the angles come in multiples, with individual indexes, the multiplicity corresponding to the announced number of photons declared [6]. The total number of photons declared is then M.
- A sequence of binary digits $\{t_k\}_k$ reported by the server, where the last digit t_M is zero.

By assumption, the server measures at least one pulse for which he gets one photon and declares one photon. Without loss of generality, let us assume that this is the case for the final pulse characterised by its polarization angle σ_M. The client will then calculate the value $\theta = \sum_{i=1}^{M-1} \sigma_i + \sigma_M$.

On the server's side, prior to declaring the binary digit outcomes, the server's quantum state can in all generality be written as:

$$\eta^{\sigma_1,\ldots,\sigma_{M-1}} \otimes |+_{\sigma_M}\rangle\langle+_{\sigma_M}|\,,$$

where the state $|+_{\sigma_M}\rangle$ is the state of the single copy declared qubit. The rest of the server's system depends on the number measurement outcomes, but can always be written in the generic form $\eta^{\sigma_1,\ldots,\sigma_{M-1}}$. In the case the server was honest, his system will be exactly of the form:

$$\left(|+_{\sigma_1}\rangle\langle+_{\sigma_1}| \otimes \ldots \otimes |+_{\sigma_{M-1}}\rangle\langle+_{\sigma_{M-1}}|\right) \otimes |+_{\sigma_M}\rangle\langle+_{\sigma_M}|\,,$$

where the states before the final state may come in multiplicities which match the declared number measurement outcomes.

Note that, whatever procedure the server may run on the system in his possession, in the spirit of the Stinespring dilation theorem, can always be represented as a unitary transform U on the input system, augmented by an ancillary system, followed by a measurement on the output of

[6] The actual number of photons will be irrelevant in the end, as the protocol assumes that only one photon is kept per pulse. However, we present this lemma in a more general form.

the overall unitary transform. The classical outcome will, in general, encode the binary digits $\{t_k\}_k$ the server has to report to the client, as in the case of no report the protocol is aborted. We emphasize that we are not assuming anything about the classical outcome - it may be a result which depends on the states the server recieved from the client, it may be chosen randomly by the server, or it may be selected. The state in the server's possession prior to measurement, ρ^{prior}, can then be viewed as the result of a CPTP map (which depends on the state $\eta^{\sigma_1,\ldots,\sigma_{M-1}}$) applied on the state $|+_{\sigma_M}\rangle\langle+_{\sigma_M}|$

$$\rho^{\text{prior}} = \mathcal{E}^{\eta^{\sigma_1,\ldots,\sigma_{M-1}}}(|+_{\sigma_M}\rangle\langle+_{\sigma_M}|).$$

This is only possible because the angle σ_M does not depend on any other angles. For simplicity, we shall fix the angles $\sigma_1,\ldots,\sigma_{M-1}$ and simply write the state prior to measurement:

$$\rho^{\text{prior}} = \mathcal{E}(|+_{\sigma_M}\rangle\langle+_{\sigma_M}|).$$

Since the angle σ_M was chosen uniformly at random, known to the client, the state of the server's system is:

$$\pi_{\text{server}} = 1/8 \sum_{\sigma_M} \mathcal{E}(|+_{\sigma_M}\rangle\langle+_{\sigma_M}|).$$

Following this, the server will measure a part of his subsystem, obtaining the sequence of binary digits $\overrightarrow{t} = \{t_k\}_k$ which he reports to the client. The state of the system after measurement (taking into account all possible outcomes) can be written as:

$$\pi_{\text{server}} = 1/8 \sum_{\sigma_M} \sum_{\overrightarrow{t}} p_{\sigma_M}(\overrightarrow{t}) \mathcal{E}_{\overrightarrow{t}}(|+_{\sigma_M}\rangle\langle+_{\sigma_M}|),$$

where $p_{\sigma_M}(\overrightarrow{t})$ is the probability of outcome $\overrightarrow{t}$ given that the input state was σ_M, and $\mathcal{E}_{\overrightarrow{t}}$ are the quantum operations which depend on the outcome. Note that:

$$\sum_{\overrightarrow{t}} p_{\sigma_M}(\overrightarrow{t}) \mathcal{E}_{\overrightarrow{t}} = \mathcal{E} \tag{2.24}$$

for all angles σ_M. As the two sums commute, we can write this state as:

$$\pi_{\text{server}} = 1/8 \sum_{\overrightarrow{t}} \sum_{\sigma_M} p_{\sigma_M}(\overrightarrow{t}) \mathcal{E}_{\overrightarrow{t}}(|+_{\sigma_M}\rangle\langle+_{\sigma_M}|). \tag{2.25}$$

Recall that a fixed sequence $\overrightarrow{t}$ along with the fixed sequence of angles $\sigma_1,\ldots,\sigma_{M-1}$ defines the angle θ:

$$\theta = \textstyle\sum_{k=1}^{M-1} (-1)^{t_k} \sigma_k + \sigma_M.$$

Note that the value of θ attains all possible angles when σ_M goes through all possible angles, for $\overrightarrow{t}$ and $\sigma_1,\ldots,\sigma_{M-1}$ fixed. Now, since the sum:

$$\sum_{\sigma_M} p_{\sigma_M}(\overrightarrow{t}) \mathcal{E}_{\overrightarrow{t}}(|+_{\sigma_M}\rangle\langle+_{\sigma_M}|)\,,$$

for a fixed sequence $\overrightarrow{t}$ goes through all the possible angles, this sum is, for every sequence $\overrightarrow{t}$, equal to:

$$\pi_{\text{server}} = 1/8 \sum_{\overrightarrow{t}} \sum_{\theta} p_\theta(\overrightarrow{t}) \mathcal{E}_{\overrightarrow{t}}(|+_\theta\rangle\langle+_\theta|).$$

Due to property (2.24), this final state is of the form $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$. □

Lemma 15. *If the server measured a weak coherent pulse sent by the client to contain zero photons, and declared it as containing one photon to the client, and the client did not abort in the presented remote blind qubit state preparation protocol, then the state shared by the client and the server after the termination of the protocol is of the form $\mathcal{E}(|+_\theta\rangle\langle+_\theta|)$ for some CPTP map $\mathcal{E}$.*

Proof. We will extensively use the setup and the arguments of the proof of Lemma 14. Assume that it is the l^{th} declared photon that the server does not possess. Then for the sequence of binary digits $(t_1 \ldots, t_k)$ the server will have reported as the alleged classical outcome of the interlaced 1D cluster computation, the angle the client computes is given as:

$$\theta = \sum_{i=1}^{k} (-1)^{t_k} \sigma_i \,,$$

which can be written as:

$$\theta = \sum_{i \in \{1,\ldots,l-1,l+1,\ldots k\}} (-1)^{t_i} \sigma_i + (-1)^{t_l} \sigma_l \,.$$

Let us fix all the σ_i angles except σ_l and all the reported binary digits t_i except t_l. The general state with the server after the remote blind state preparation protocol can then be written as:

$$\pi_{\text{server}} = \sum_{t_l=0}^{1} p(t_l) \sum_{\sigma_l} \frac{1}{8} \eta^{t_l}$$

where η is the state in the hands of the server, which may depend on t_l, and $p(t_l)$ is the probability of the server reporting t_l to be one or zero. Note that neither the probability $p(t_l)$ nor the final state η^{t_l} can depend on σ_l. The angle θ in the expression above for any fixed t_l goes across all possible values as σ_l ranges across all possible values. Hence, the sum may be written in terms of the angle θ rather than σ_l as:

$$\pi_{\text{server}} = \sum_{t_l=0}^{1} p(t_l) \sum_{\theta} \frac{1}{8} \eta^{t_l} \,.$$

Also, since η^{t_l} does not depend on σ_l it does not depend on θ so we can factor it out of the sum

$$\pi_{\text{server}} = \sum_{t_l=0}^{1} p(t_l) \left(\sum_{\theta} \frac{1}{8} \right) \eta^{t_l} = \sum_{t_l=0}^{1} p(t_l) \eta^{t_l} \,.$$

Let η be the quantum state $\sum_{t_l=0}^{1} p(t_l)\eta^{t_l}$, then we have

$$\pi_{\text{server}} = \eta$$

which implies

$$\pi_{\text{server}} = \mathcal{E}^{\eta}(|+_\theta\rangle\langle+_\theta|)$$

where $\mathcal{E}^{\eta}$ is a CPTP map which is the contraction to the fixed state η. □

This lemma proved that the server cannot gain anything from declaring one (or more) photons when he in fact received none (which is one way to attempt to cheat, and decrease the number of reported vacuums). Thus, we could eliminate this type of malicious behaviour from the analysis.

2.7 Discussion

As presented the RBSP protocol is not immune to noise in the channel or to significant preparation errors on the side of the client. In particular, throughout the run of the I1DC subroutine, explained in Protocol 5, one can see that whatever initial preparation errors may have, or which may have occurred in the transmission, accumulate in the resulting "blind" qubit the server will use in UBQC. This means that in order to satisfy a threshold η the fault tolerant code can handle, the states the client delivers to the server in RBSP have to satisfy the threshold η/N, where N is the security parameter of RBSP. While N for the required level of security scales only logarithmically, N may still in the end be large for large computations, in which case we actually demand that the client's preparation be much better than the precision of the servers devices. This is not completely satisfactory, but given that the client only needs to do this one particular task well (rather than a full blown quantum computation), perhaps it could still be made to work. The solution would be to find a method of performing RBSP in a fault tolerant way. One means of doing this would be by adapting techniques used to ensure the fault tolerance of UBQC itself [54, 1, 47]. Additionally, we could increase the number of parameters of the protocol and allow the client to control the norm of the amplitude (mean photon number) of the pulses as well. This will give the client more information with which she could track the server's behaviour better. Such an approach was successfully used in QKD with coherent states, and is often called QKD *with decoy pulses* [62, 63]. Perhaps similar schemes could be used in our case, without making the server significantly more complicated. The potential solutions to this problem we leave for future research. However, we emphasize that noise can only jeopardize the correctness of our protocol, but never the guaranteed security levels.

So far we have considered Alice's side of the story. In what follows we address the side of Bob. We consider generalized MBQC, and begin an investigation into whether and how UBQC can be performed on a modified AKLT state, which has advantages in the sense of resilience to decoherence.

Chapter 3

The world according to Bob: *UBQC using alternative resources*

We investigate the feasibility of universal blind quantum computation based on AKLT underlying resource states, and consider the interplay between "ground state-ness" of a blind computation, imposed communication assumptions and cryptographic privacy

3.1 Robustness of the measurement-based computation models

In the last chapter we addressed the question of the feasibility of UBQC in *practice* with all the unfortunate imperfections reality brings, from the perspective of the client. We have shown that the client's privacy, that is, blindness can be maintained even with realistic assumptions on imperfections. We have noted that only the imperfections on the side of the client can jeopardize blindness. However, *the correctness* of the protocol depends crucially on the side of the server – he has to have a working robust and scalable quantum computer. The problem of *robust and scalable* quantum computation is arguably *the problem* in quantum computation for a large part of the scientific community. In this chapter we do not attempt to contribute to this problem, but rather attempt to apply the existent proposals to the cryptographic setting of UBQC. Already in the original proposal [1] it was noted that certain levels of fault tolerance can be achieved in UBQC by using known fault tolerant constructions in quantum circuits. Recently, more explicit constructions of fault tolerant codes, in particular the Raussendorf-Goyal-Harrington [47] topological code based in the qubit-based MBQC have been adapted for UBQC [54], and even thresholds were computed. However, all of these solutions still assume that the server runs the computation on a qubit-based graph state. However, more robust computation can also be achieved by considering different physical systems which are intrinsically more resilient to the main "killer" of quantum computation – decoherence. The recent interaction between condensed matter physics and quantum information resulted in novel measurement-based models of computations, which use resource states which may have advantages, concerning the robustness of the computation, over the well-studied graph states. In particular, the resource state we will consider is a gapped ground state of a natural Hamiltonian, meaning it could in principle be prepared by cooling, and that the computation itself could be protected from decoherence by keeping the quantum computer cold enough.

In this chapter we introduce the basics of these novel ideas and address the question of how

UBQC can be adapted to work with these recent, perhaps more experimentally friendly, computational models. In the process we touch upon an a potentially interesting interplay between central notions in condensed matter physics, cryptography and classical communication.

3.2 Matrix Product States and Generalized MBQC

"The Hilbert space is a big place" – Carlton M. Caves. Exponentially big, to be exact: to uniquely fix a general pure state of a system comprising n d-dimensional subsystems one needs to specify roughly[1] d^n complex parameters. The intractability of even writing down the state of such a system, for any non-minuscule choice of n prompted the research of states which do allow for an efficient representation, and are physically interesting. A very successful family of such states that we have extensively used in the presented work were graph states which succinctly described the resource states of n qubit systems used in MBQC. Another stellar example are the more general stabilizer states [10]. As mentioned, stabilizer states are ubiquitous in the game of error-correcting codes and fault-tolerant quantum computation, and were also used in the proof fo the fabled Gottesman-Knill theorem, which characterised a very exciting family of quantum computations that can be classically simulated. We have also used the stabilizer formalism and a particular family of stabilizer states – graph states, to perform deterministic one-way quantum computation discussed in Chapter 1.

Matrix product representation of quantum states, in its most general form allows for the representation of all quantum states. Let $|\psi\rangle \in \bigotimes_{i=1}^{n} H^d$ be any general state of an N-partite system comprising d-dimensional subsystems. Then, we can represent this state with the expression

$$|\psi\rangle = \sum_{(l_N,\ldots l_1)=(0,\ldots,0)}^{(d-1,\ldots,d-1)} c(l_N,\ldots,l_1)|l_N\rangle|l_{N-1}\rangle\ldots|l_1\rangle, \tag{3.1}$$

where the states $\{|l_k\rangle\}_{k=0}^{d-1}$ comprise an orthonormal basis of the individual subsystem state space H^d and the coefficients $c(b_n,\ldots,b_1)$ are complex numbers. In a matrix product representation of the same state the coefficients $c(l_N,\ldots,l_1)$ assume a particular form:

$$c(l_N,\ldots,l_1) = \overrightarrow{L}^T.A_N\left[l_N\right]\cdots A_1\left[l_1\right].\overrightarrow{R}, \tag{3.2}$$

where L and R are D dimensional numerical vectors, and $A_i[b_k]$ is a $D\times D$ complex matrix for each $k=0,\ldots,d-1$. Note that, in the right-hand side of the expression above, what we have is the standard inner product between the vector $\overrightarrow{L}$ and the vector $A\left[l_N\right]\cdots A\left[l_1\right].\overrightarrow{R}$. By abuse of notation, which as we shall see is highly suggestive, we shall write the expression above in the bra-ket formalism as

$$c(l_N,\ldots,l_1) = \langle L|A_N\left[l_N\right]\cdots A_1\left[l_1\right]|R\rangle. \tag{3.3}$$

[1] One may argue that we need less than this, given the restriction that the sum of squared norms of these parameters needs to be unity.

In general D and d may be unrelated, and clearly by selecting D large enough (it may depend on n), any state can be represented in this form [64]. However, here we will be interested in states in which the matrix product representation is efficient. We will require D to be independent of n and small, and that $A_i[l_k] = A_j[l_k]$. In particular we require $D = 2$. The set of states which are representable in this form plays an important role in condensed-matter physics [65, 66], and we shall give two examples of such relevant states but applied to the needs quantum information processing. Before this we first introduce the concept of correlation space [48, 50, 51].

Consider the matrix-product state of a quantum state in the d^N-dimensional Hilbert space

$$\sum_{l_1=0}^{d-1} ... \sum_{l_N=0}^{d-1} \langle L|A[l_N]...A[l_1]|R\rangle |l_N\rangle \otimes ... \otimes |l_1\rangle,$$

where $|L\rangle$ and $|R\rangle$ are two-dimensional complex vectors and $A[0], \ldots, A[d-1]$ are two-dimensional complex matrices. Let us assume that the first qudit of the matrix-product state is projected onto

$$|\theta, \phi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle,$$

as a result of a particular measurement. Then, the post measurement state can be represented in the following matrix-product-like form:

$$\sum_{l_2=0}^{d-1} ... \sum_{l_N=0}^{d-1} \langle L|A[l_N]...A[l_2]A[\theta, \phi]|R\rangle |l_N\rangle \otimes ... \otimes |l_2\rangle \otimes |\theta, \phi\rangle,$$

where

$$A[\theta, \phi] = \cos\frac{\theta}{2}A[0] + e^{-i\phi}\sin\frac{\theta}{2}A[1].$$

If $A[0]$ and $A[1]$ are appropriately chosen in such a way that $A[\theta, \phi]$ is unitary, we can "simulate" a unitary rotation $A[\theta, \phi]$ of $|R\rangle$ in the linear space where $|L\rangle$, $|R\rangle$, and A's live. This linear space is called "correlation space". In framework of generalized measurement-based quantum computation, (often called the computational tensor networks [48, 67, 68]), universal quantum computation is performed in this correlation space.

This separation between the correlation space and the physical space allows us to use many new resource states for measurement-based quantum computing.

3.2.0.2 Example: Graph state MBQC as MPS state generalized MBQC

A simple example of such a computation we have already encountered in terms of the MBQC on graph states, as presented in [67].

Consider a graph state $|G\rangle$, which is simply a path of N qubits. It is easy to verify that this is

also a MPS state of the following structure:

$$|G\rangle = \sum_{(l_N,\ldots,l_1)=(0,\ldots 0)}^{(1,\ldots,1)} \langle 0|A\left[l_N\right]\cdots A\left[l_1\right]|+\rangle |l_N \ldots l_1\rangle, \tag{3.4}$$

where

$$A[0] = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} \text{ and } A[1] = \begin{bmatrix} 0 & 1 \\ 0 & -1 \end{bmatrix}. \tag{3.5}$$

A measurement of the first qubit with respect to the observable M^{ϕ_1} (which collapses the qubit state to one of the states $|\pm_{\phi_1}\rangle$ as described in Section 1.2.1) in the event of the outcome $s = 0$ corresponding to the $|+_{\phi_1}\rangle$ post-measurement state modifies the state $|G\rangle$ to

$$|G^1\rangle = \sum_{(l_N,\ldots,l_2)=(0,\ldots 0)}^{(1,\ldots,1)} \langle 0|A\left[l_N\right]\cdots A\left[l_2\right]\left(A\left[0\right] + e^{-i\phi_1}A\left[1\right]\right)|+\rangle |l_N \ldots l_2\rangle |+_{\phi_1}\rangle, \tag{3.6}$$

and, not surprisingly $A\left[0\right] + e^{-i\phi}A\left[1\right] = HZ_{-\phi_1}$. Recall, the latter parametrized unitary we denoted as the "J" gate: $J(\phi) := HZ_\phi$. Repeating this procedure, we obtain the state

$$|G^{N-1}\rangle = \sum_{(l_N)=(0)}^{(1)} \langle 0|A\left[l_N\right]J(-\phi_{N-1})\cdots J(-\phi_1)|+\rangle |l_N\rangle |+_{\phi_{N-1}}\rangle \cdots |+_\phi\rangle, \tag{3.7}$$

As we can see, what we have done by this procedure is compute a sequence of J gates in the correlation space, which is universal for single-qubit computation. The probabilities of the measurement outcome of the last non-measured qubit with respect to the Z observable are equal to the probabilities of the measurement of the state $J(-\phi_{N-1})\cdots J(-\phi_1)|+\rangle$, which lives in the correlation space, with respect to the same observable.

For the case of 1D graph states, as we have just demonstrated, the correlation space and the physical space coincide – we have just recaptured exactly the same results as in Section 1.2.1 which explicitly dealt with the measurement-induced *physical* state evolution of graph states. However, for the case of the next example of an MPS state which we will be working with for the remainder of this chapter, this will not be the case. One such resource state is the AKLT state, named after Affleck, Kennedy, Lieb, and Tasaki [65].

3.2.1 The AKLT state

In this section we quickly define the AKLT state and its basic properties. Following this, we will show it is a universal resource for quantum computation, and therefore a viable candidate for a new resource in UBQC.

We consider a one-dimensional open-boundary chain of N spin-1 particles (*i.e.* qutrits, $d = 3$).

The AKLT Hamiltonian [65] is defined by

$$H_{AKLT}(\beta) \equiv \sum_{j=1}^{N-1} h_{j+1,j}(\beta),$$

where

$$h_{j+1,j}(\beta) \equiv \frac{1}{2}[\mathbf{S}_{j+1} \cdot \mathbf{S}_j - \beta(\mathbf{S}_{j+1} \cdot \mathbf{S}_j)^2]$$

and $\mathbf{S}_j \equiv (S_j^x, S_j^y, S_j^z)$ is the spin-1 operator on site j defined by

$$\begin{aligned} S_j^x &\equiv \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \\ S_j^y &\equiv \frac{-i}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix}, \\ S_j^z &\equiv \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \end{aligned}$$

If $-1 < \beta < 1$, the system is in the so-called gapped Haldane phase [69]. For $\beta = -1/3$, the ground states are called AKLT states [65] and explicitly written in the following matrix product form [48, 66]:

$$|AKLT^{N,L,R}\rangle \equiv \frac{\sqrt{2}}{\sqrt{3^N}} \sum_{l_1=1}^{3} ... \sum_{l_N=1}^{3} \langle L|A[l_N]...A[l_1]|R\rangle |l_N\rangle \otimes ... \otimes |l_1\rangle, \tag{3.8}$$

where

$$\begin{aligned} |1\rangle &\equiv -\frac{1}{\sqrt{2}}\Big(|S_z = +1\rangle - |S_z = -1\rangle\Big), \\ |2\rangle &\equiv \frac{1}{\sqrt{2}}\Big(|S_z = +1\rangle + |S_z = -1\rangle\Big), \\ |3\rangle &\equiv |S_z = 0\rangle, \end{aligned}$$

$|S_z = k\rangle$ ($k \in \{-1, 0, +1\}$) are eigenvectors of the z-component S_z of the spin-1 operator, $|L\rangle$ and $|R\rangle$ are two-dimensional complex vectors, and $\{A[1], A[2], A[3]\}$ are 2×2 matrices defined by

$$\begin{aligned} A[1] &\equiv X, \\ A[2] &\equiv XZ, \\ A[3] &\equiv Z. \end{aligned}$$

As we can see, in the example of the AKLT state, the dimensionality of the realized correlation space does not match the dimensionality of the underlying resource subsystems. The Pauli matrices above are qubit (2 dimensional) matrices, and so are the vectors $|L\rangle$ and $|R\rangle$, whereas the underlying resource is built up from (3 dimensional) qutrits - spin 1 particles. The ground states of the AKLT Hamiltonian are four-fold degenerate, and each ground state is specified with the choices of $|L\rangle$ and $|R\rangle$. The AKLT states are (the ground states of a) frustration-free (2-local nearest neighbour Hamiltonian), since

$$h_{j+1,j}(-1/3)|AKLT^{N,L,R}\rangle = 0$$

for any $|L\rangle$ and $|R\rangle$, and for all $j = 1, ..., N-1$.

Recall, a *local* Hamiltonian $H = \sum_j H_j$ is frustration-free if H_j are positive semi-definite operators and the ground-state of H is a zero eigenvector of all operators H_j. Locality, more precisely k-locality here implies that every term H_j of the Hamiltonian H acts non-trivially on k or fewer qudits (qutrits, in our case).

Such frustration-free 2-local nearest neighbour Hamiltonians are of interest as they are considered *natural* – meaning they may occur in nature, and more importantly for our purposes, could potentially be prepared in a lab. The AKLT states are of immediate interest as they are the ground states of such a Hamiltonian, and could in principle be prepared by cooling.

The AKLT model has a long history in condensed matter physics. It describes the one-dimensional Haldane phase [69] of a qutrit chain, which, for instance, exhibits the system size-independent spectral gap [65] and the effective spin-1/2 degree of freedom, namely the edge state, appearing on the boundary of the chain [70] (again, represented in terms of the "boundary conditions" $|L\rangle$ and $|R\rangle$). Furthermore, the AKLT model has recently been attracting increasing attention in quantum information, because of its connections to the matrix product representation [66, 71], localizable entanglement, and generalized measurement-based quantum computation [48, 67]. Indeed, it was shown in [49] that the measurement-based quantum computation is possible on the AKLT chains and other ground states in the gapped Haldane phase ($-1 < \beta < 1$). Next we describe the basic elements of this result.

3.2.1.1 MBQC on AKLT

We briefly review universal measurement-based quantum computation on AKLT states [49]. Assume that a qutrit of the system represented in expression3.8 is measured in the basis $\mathcal{M}(\phi) = \{|\alpha(\phi)\rangle, |\beta(\phi)\rangle, |\gamma\rangle\}$, where

$$\begin{aligned} |\alpha(\phi)\rangle &= \frac{1+e^{i\phi}}{2}|1\rangle + \frac{1-e^{i\phi}}{2}|2\rangle, \\ |\beta(\phi)\rangle &= \frac{1-e^{i\phi}}{2}|1\rangle + \frac{1+e^{i\phi}}{2}|2\rangle, \\ |\gamma\rangle &= |3\rangle. \end{aligned}$$

Then, following operations are implemented in the correlation space according to the measurement result [49].

$$\begin{aligned} |\alpha(\phi)\rangle &: \quad XZ_\phi, \\ |\beta(\phi)\rangle &: \quad XZZ_\phi, \\ |\gamma\rangle &: \quad Z, \end{aligned}$$

where $Z_\phi = |0\rangle\langle 0| + e^{i\phi}|1\rangle\langle 1|$ (equivalently $Z_\phi = e^{-i\phi Z/2}$, up to an irrelevant global phase). We will use analogous notation for Pauli X, so $X_\phi = e^{-i\phi X/2}$. If the unitary operation

$$V = |3\rangle\langle 1| + |1\rangle\langle 2| + |2\rangle\langle 3|$$

is applied on a qutrit and that qutrit is measured in the basis $\mathcal{M}(\phi)$. Then, following operations are implemented in the correlation space according to the measurement result (represented by the post-measurement state) [49].

$$\begin{aligned} |\alpha(\phi)\rangle &: \quad XZX_\phi, \\ |\beta(\phi)\rangle &: \quad ZX_\phi, \\ |\gamma\rangle &: \quad X. \end{aligned}$$

As we can see, in the two cases above, where the unitary V was not or was applied, Pauli Z eigenvalue axis and Pauli$X$ eigenvalue axis rotations of the Bloch sphere, respectively, are implemented in the correlation space, provided outcomes $|\alpha(\phi)\rangle$ or $|\beta(\phi)\rangle$ were obtained up to Pauli byproducts. By appropriately modifying the sign of ϕ in subsequent measurements (adapting measurements of the physical qutrits, much like in regular MBQC presented earlier), these byproducts can be shifted "forward" and corrected in the final stage of the computation over the correlation space. In particular, this correction can be done on the classical final measurement outcome. Note that the computation in this model is not deterministic. With a non-zero probability, each measurement may result in the outcome state $|\gamma\rangle$ (for both axes), in which case just a byproduct is implemented. Then, the measurement needs to be repeated whenever this undesirable "third outcome" is obtained.

This enables us to perform any single qubit quantum computation on the AKLT resource. To achieve computational universality, a two qubit interaction is required. This can be done by applying a two-qutrit interaction on the qutrits of two AKLT chains, followed by local measurements of the two qutrits. Thus, in this particular model of generalized MBQC, unlike in the graph state, the entanglement of "computational lines" (see Figure 1.4) is performed in run time, rather than in preparation of the resource. The success probability of the application of this entanglement operation is again not unity, and the process may have to be repeated multiple times until it succeeds. This we explain in our UBQC protocol over the AKLT states in Section 3.3.

We note that MBQC on an AKLT state was recently experimentally demonstrated in an optical systems [72].

3.3 UBQC with AKLT

As we have seen, similarly to the previously discussed graph-states in the one-way model, the AKLT states are a universal resource for generalized MBQC. The AKLT states, however, hold certain physical advantages over graph states. Namely, they are ground states of natural Hamiltonians, and this property could be used to reduce external noise introduced to the system during quantum computation by keeping the system cool. In contrast, graph states are known not to be ground states of natural Hamiltonians. While this advantage has been exploited for quantum computation, here we investigate whether it can also be used in UBQC. An affirmative result could help make Bob's side of the story more realistic. We begin by very briefly outlining the original graph-state based UBQC protocol in order to isolate the key moments we wish to capture in an attempt to construct an AKLT-based version of UBQC.

3.3.1 UBQC on graph states

Let us assume that Alice, the client, wants to perform the MBQC on an N-qubit graph state $|G\rangle$ with the measurement on jth qubit in the basis $\{Z_{\phi_j}|+\rangle, Z_{\phi_j}|-\rangle\}$. If Alice asks Bob, the server, to create $|G\rangle$ and sends $\{\phi_j\}_{j=1}^N$ to Bob, they can perform the correct delegated computation. However, in this case, Bob can learn Alice's computation. To prevent this, Alice "encrypts" the computational angles as follows: Alice first sends randomly rotated single-qubit states, $\{Z_{\theta_j}|+\rangle\}_{j=1}^N$, to Bob, where θ_j works as a part of the "encryption key". Bob applies CZ gates among them, as dictated by the generic resource graph. Since Z_{θ_j} commutes with CZ gates, what Bob has is $\bigotimes_{j=1}^N Z_{\theta_j}|G\rangle$, where Z_{θ_j} is acting on jth qubit. If Alice sends $\phi_j + \theta_j + r_j\pi$ to Bob, where $r_j \in \{0, 1\}$ is chosen uniformly at random, Alice's true computational angle ϕ_j is one-time padded with $\theta_j + r_j\pi$, and hence Alice can have Bob do the correct MBQC without ever revealing ϕ_j. For more details on significance of the additional hidden "r" parameter, see Section 4.3.2.

Characteristic elements of UBQC The UBQC protocol ensures that the computation the server performs for the client remains hidden from the server. Since the server *does* perform the desired computation, and yet learns nothing about it, this protocol must employ some type of encryption. In the case of UBQC it effectively the *hardware* that is encrypted, by encryption of the resource state. The encryption used in UBQC stems from a few very convenient properties of the one-way model. In the one-way model, the required measurements for driving the computation are single qubit measurements parametrized by an angle ϕ the observables of which we denoted M^ϕ. Note that, for $\phi = 0$ this observable is simply the Pauli X, and more generally it holds that $M^\phi = Z_\phi X Z_{-\phi}$, for $Z_\phi = |0\rangle\langle 0| + e^{i\phi}|1\rangle\langle 1|$. Thus, a Pauli-$Z$ pre-rotation of an individual qubit in the resource graph state corresponds to the modification of the effective measurement angle. Because of this reason, the encryption of the resource state by local pre-rotations can completely conceal the *semantics* of the physical measurement performed. That is, conceal the angles defining the desired computation in the positive branch performed over an unencrypted

resource state, and hide the intended computation. However, if in order to make UBQC feasible we would have to require the client generate such a pre-rotated resource state UBQC would not be very practical for the client. Here, the other convenient property of the one-way model comes to the rescue – the required individual qubit pre-rotations commute with the entanglement operation. Thus, to ensure the server has an encrypted resource state, the client only needs to prepare individual pre-rotated qubits, and it will be the servers job to entangle them according to a pre-defined universal and generic graph.

To summarize, the two key aspects which make UBQC feasible based on the one-way model then are: 1) That a "pre-rotation" of the sub-systems of the resource state which will be individually measured can conceal the intended measurement, and 2) That the process of generating such a large encrypted resource state can be done by first generating a type of encrypted small systems (something a client can realize using the simplest devices possible) which are then transformed to the final encrypted resource state by means of a fixed interaction. Equally importantly, in this chapter we are interested in simplifying the already formidable requirements on the server. So, the interaction the server needs to perform to generate the large encrypted resource state from the smaller "pre-rotated" components, the fixed interaction should be as simple as possible as well.

Property 1) seemingly can clearly always be ensured - any parametrized measurement can be realized by first applying a unitary (parametrized by the measurement parameter) followed by a fixed measurement (this is exactly what unitary transforms do – change the local basis). While this observation may seem to resolve the first issue in MBQC based protocols, in the case of AKLT computations we may need to exercise a bit of caution. Unlike in the case of the one-way model, as we have seen in the AKLT case the computation is not happening in the physical space of the subsystems, but rather in the correlation space. How unitaries (or general CPTP maps) affect the virtual states in the correlation space is not straightforward. For instance, it has been shown that CPTP errors in the physical space may correspond to non-trace-preserving disturbances in the correlation space, which makes standard fault-tolerant techniques not applicable in generalized MBQC models [73]. Note that while the correlation space certainly is linear, it definitely need not manifest unitary maps only. Surely, if UBQC is to be realized in the AKLT setting these problems need to be accounted for.

Perhaps even more seriously, property 2) seems to be specific to the one-way model. A direct naive attempt to generalize how this is resolved in the original UBQC protocol would work as follows: The client sends somehow "pre-rotated" qutrits. The server builds up the encrypted resource state from them using fixed local interactions. Unfortunately, this is not possible. It is known that the AKLT state cannot be created from separable states with local (k-local, where k does not depend on the size of the desired final AKLT state) unitary operations [53] It is true that, in principle, Bob possibly could rebuild some type of an encrypted ALKT state from pre-rotated qutrits, using elaborate and fixed [2] global operations. It is an open question whether such a global process, which would work with sensibly and result with an encrypted state useful in the sense of property 1), exists. However, even if it did, it would still require the server to

[2] By fixed we mean independent from the pre-rotation angles.

perform global operations which are arguably a lot more demanding than what he needs to do in the original UBQC protocol – in an attempt to help poor Bob, we would have made his life even more difficult! For this reason, we will dismiss this direct attempt. This issue will be resolved using a different approach.

Finally, even if all these problems are resolved, there is still one remaining fundamental difference between the one-way model and AKLT-based generalized MBQC. As we have seen, in the ideal case, one-way computation is deterministic. However, for each measurement we do in the AKLT case, there is a non-negligible probability that the measurement returns the "third outcome" (see equation 3.9) in which case just a Pauli byproduct was generated instead of the desired measurement angle determined rotation. Thus AKLT computation is probabilistic. To combat this problem, in the case of this trivial measurement outcome, the measurement needs to be repeated. This opens the door for a plethora of malicious activities for the server as he, in the case of this trivial outcome, *knows* the following computational measurement angle will have to be a repetition of the prior measurement angle. Such additional information was not accessible to the server, so if the AKLT based UBQC is at all possible, it will require a new and more elaborate security proof.

Now we proceed to show how all of these problems are resolved, but at a substantial cost which highlights an interesting interplay between the cryptographic notions of privacy and central concepts in condensed-matter physics.

3.3.2 A UBQC protocol with AKLT states

As we mentioned, to have a satisfactory AKLT-based UBQC protocol, it should satisfy the following three properties: 1) The encrypted resource state the server builds up from the components the client sends should conceal the computation the client wishes to run *in the correlation space*, 2) the encrypted resource state should be built up from simple elements using a simple interaction and 3) the entire protocol should be secure even in the face of the probabilistic nature of generalized MBQC. We first focus on the second property, which as we have mentioned cannot be resolved by using any type of "pre-rotated" qutrits. We will show how building a type of an encrypted resource state is possible using PEPS (projected entangled pair states) projections, a key idea originally noticed by Tomoyuki Morimae. From the construction it will be apparent, at the intuitive level at least, that this type of encryption conceals the computational angles in the same way as in the original protocol. Thus, property 1) is satisfied as well.

Then, due to its technical nature, in a separate Section 3.6.1 we will give a detailed formal proof of security of the proposed single server AKLT UBQC protocol which takes into account the probabilistic nature of AKLT-based computation. As is shown in Figure 3.1, Alice has a classical computer (C Comp) and a quantum device (QD) which emits random four-qubit states

$$(I \otimes (I - |\eta_1\rangle\langle\eta_1|) \otimes I)(I \otimes T_{Z/X}(\theta) \otimes I)|\eta_1\rangle \otimes |\eta_1\rangle,$$

where $\theta \in \{\frac{k\pi}{4}|k = 0, ..., 7\}$ is a random angle, $|\eta_1\rangle \equiv (|00\rangle + |11\rangle)/\sqrt{2}$, $|\eta_2\rangle \equiv (I \otimes Z)|\eta_1\rangle$,

$|\eta_3\rangle \equiv (I \otimes X)|\eta_1\rangle$, $|\eta_4\rangle \equiv (I \otimes XZ)|\eta_1\rangle$, $T_Z(\theta) \equiv |00\rangle\langle 00| + e^{i\theta}|01\rangle\langle 01| + |10\rangle\langle 10| + |11\rangle\langle 11|$, and $T_X(\theta) \equiv (\cos\frac{\theta}{2}|\eta_2\rangle - i\sin\frac{\theta}{2}|\eta_4\rangle)\langle\eta_2| + |\eta_1\rangle\langle\eta_1| + (-i\sin\frac{\theta}{2}|\eta_2\rangle + \cos\frac{\theta}{2}|\eta_4\rangle)\langle\eta_4| + |\eta_3\rangle\langle\eta_3|$. Each of them is directly sent to Bob through the quantum channel (QC) (Figure 3.1 (a)). These pairs of bell states will fulfil the purpose of the building blocks of the final encrypted resource state, analogous to single qubits in the original UBQC scheme. The two qubit pre-rotation $T_{Z/X}(\theta)$ plays an analogous role to the qubit pre-rotation angle in the original UBQC. Note that in the AKLT computation rotations with respect to Pauli-X axis and Pauli-Z axis need to be performed separately. The pre-rotation incurred on the building blocks $T_{Z/X}(\theta)$ will also have to depend on which rotation we wish to implement. Any computation can be designed in a generic way, as a sequence of alternating X and Z rotations, so this alone does not jeopardize the security. For more details see Section 3.6.1.

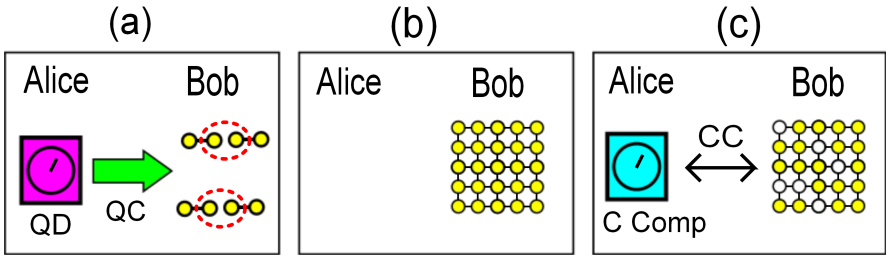

Figure 3.1. The single-server protocol. In (a), two (yellow) circles connected by a solid bond represent $|\eta_1\rangle$. The operation $(I - |\eta_1\rangle\langle\eta_1|)T_{Z/X}(\theta)$ is applied on two qubits specified with dotted (red) circles.

From these states, the server creates the "encrypted" resource state using PEPS projections as explained in what follows and Figure 3.1.

Building the encrypted resource state Bob starts from Figure 3.2 (a). Bob applies the filtering $\{|\eta_1\rangle\langle\eta_1|, I - |\eta_1\rangle\langle\eta_1|\}$ to each pair of two qubits specified by a dotted blue circle in Figure 3.2 (b). If $|\eta_1\rangle\langle\eta_1|$ is realised, two qubits are just removed from the chain (Figure 3.2 (c)). Bob next applies PEPS operation

$$P \equiv \frac{1}{\sqrt{2}} \sum_{l=1}^{3} \sum_{i=0}^{1} \sum_{j=0}^{1} A_{i,j}[l]|l\rangle\langle i| \otimes \langle j|$$

on every pair of qubits (Figure 3.2 (d)). It can be shown that $PT_Z(\theta) = U(\theta)P$ and $PT_X(\theta) = V^\dagger U(\theta)VP$. All PEPS operations are done deterministically since $I - |\eta_1\rangle\langle\eta_1|$ is applied to every pair of qubits before the PEPS operation. Bob now has a one-dimensional chain of qutrits where each qutrit is randomly rotated by U or $V^\dagger UV$ (Figure 3.2 (f)). We call such a chain an "encrypted AKLT state". A single chain of the encrypted AKLT state is used for the single-qubit rotation in the correlation space. Recall, in the graph state UBQC case, the resouce state is encrypted by rotating each qubit by $e^{i\theta Z/2}$. Here, in the AKLT case, $U(\theta)$ and $V^\dagger U(\theta)V$ correspond to the "encryption" operation $e^{i\theta Z/2}$. $U(\theta)$ is the encryption operation for Z rotations, and $V^\dagger U(\theta)V$ is that for X rotations.

Once the "encrypted" resource state is generated with the server, delegated computation is run

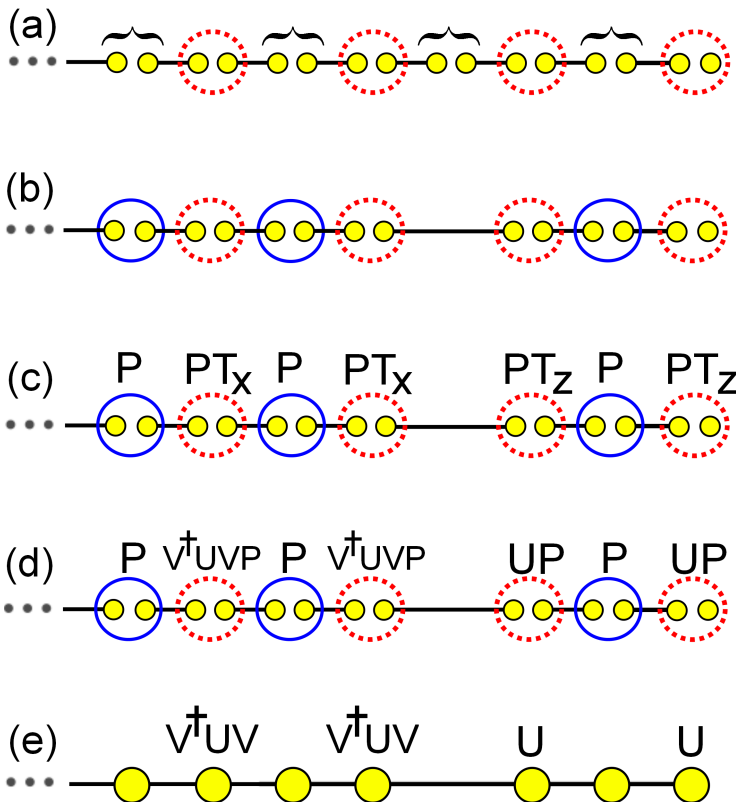

Figure 3.2. (a): $(I - |\eta_1\rangle\langle\eta_1|)T_{Z/X}$ is applied on every dotted (red) circle by Alice. Bob applies the filtering operator to every pair of qubits indicated by curly brackets. (b): If the undesired filtering outcome is realized the corresponding qubits are removed, e.g., the seventh and eighth qubits. Otherwise the qubits are entangled presented with solid (blue) circle. (c): The PEPS P is applied to every pair of two qubits. Simple rewriting leads to equivalent states (c), (d), and (e) where in the last one the large (yellow) circles are qutrits.

using two-way classical communication in an exactly the same way as in the original UBQC [1] as is shown in Figure 3.1 (c). That is, if Alice's computational angle is ϕ_j, she sends Bob the "encrypted" angle $\phi_j + \theta_j + r_j\pi$, and Bob does the measurement in that angle.

What we have described here is a very simple description of the ideas behind the AKLT-based UBQC single server protocol. The full description of the protocol, which deals with the probabilistic nature of measurements, undesirable outcomes of the filtering operations, and other details, we refer the interested reader to Section 3.6.1 For this protocol we can show that the (malevolent) server learns nothing about Alice's input, output, and algorithm despite the different resource state and the probabilistic nature of the computation by introducing new security proof techniques.

3.3.3 Trade-off between the security and the energy-gap protection

We have seen that it is possible to achieve a version of UBQC where the server performs a computation over an encrypted AKLT resource state. Recall, however, the initial motivation for this endeavour: AKLT-based computation has an advantage over a graph-state-based computation because of the gap-energy protection against noise which exists when one works with the AKLT states. In order to ensure blindness, the state the server has to work with is a modified variant of the AKLT state – it has been encrypted using a secret key, which are the pre-rotation angles. The question is, can all these states (for all possible keys which must remain hidden from the server) be ground states of a single Hamiltonian which can actually be realized? The answer to this is no. Consider the span

$$\left\{ |RAKLT_b^{2N,L,R}(\{\theta_{a,b}^{Z/X}\})\rangle \;\middle|\; \theta_{a,b}^{Z/X} \in \mathcal{A},\ a = 1, ..., N \right\}$$

of all rotated AKLT states. It can be shown that dimension of the span above is 2^{2N} (for a proof see 3.6.5).

This means that the computation the server needs to run cannot be energy-gap protected unless the ground space is exponentially degenerated, and this property is not satisfied by what are considered natural Hamiltonians. Hence we are confronted with a trade-off: in the protocol we present we cannot satisfy security and energy-gap protection simultaneously. We note that such a trade-off between the security and the energy-gap protection is not specific to the AKLT model. This trade-off holds for general resource states which are parametrized by an exponential number of keys, and are physically distinguishable, for differing choices of keys [3]. Hence, all cryptographic schemes which use the "encryption of the computation resource" and aim to enjoy the energy-gap protection require unnatural Hamiltonians (in terms of the exponential degeneracy). . While this constitutes a set-back, it also introduces an interesting interplay between central concepts in condensed-matter physics, which to our knowledge has not previously been addressed. However, a resolution of the dichotomy above is possible, by introducing a particu-

[3]For the encoding to make any sense, the states must differ for different keys in a physically meaningful way. Adding a parametrized global phase to a fixed state, for instance, will not violate gap-energy protection, but can also not be exploited in any physical process.

lar type of no-communication assumption. We now turn our attention to a two-server setting in which two central goals get resolved simultaneously: energy-gap protected computation by the servers, and a completely classical client.

3.4 Two-server UBQC with AKLT

Already in the original UBQC paper [1] a setting with two non-communicating servers, who do share entangled pairs was considered. Such interactive non-communicating multi-server computation protocols arise naturally as a generalization of interactive proof systems called multi-prover interactive proof systems, a well-studied topic in classical and quantum complexity theory [74, 75]. The authors of the original paper have already noted that, if Alice has access to two quantum servers who share entanglement, but do not communicate, then UBQC is possible even when Alice has only the powers to communicate classically (for more on this topic and the relationship between UBQC and interactive proof systems see Chapter 4). Surprisingly, the two-server setting offers additional advantages when UBQC is performed using AKLT states: If we consider two servers, Bob_1 and Bob_2, who are prohibited from communicating with each other, then both the security and the energy-gap protection can be realized! In this section we show that the energy-gap-protected UBQC is possible with the usual (four-fold degenerate) AKLT Hamiltonian in such a two-server situation. Then, in such a two-server protocol both energy-gap protection is achieved and the client can be fully classical.

We briefly recap the basic idea in the two-server scheme presented in the original UBQC paper. First, Bob_1 creates randomly-rotated single-qubit states $\{Z_{\theta_j}|+\rangle\}_{j=1}^{N}$, teleports them to Bob_2. Bob_1 reports $\{\theta_j\}_{j=1}^{N}$ and the results of the teleportations to Alice. Second, Bob_2 creates the graph resource state from the teleported qubits. Third, Alice and Bob_2 effectively perform the single-server UBQC. As we have already mentioned in [1] it was shown that Alice's privacy is ensured against both Bobs. Furthermore, in the two server setting, Alice no longer requires any quantum powers. The key idea in this protocol is that Bob_1 cannot learn Alice's computational angles since Bob_2 does measurements, whereas Bob_2 cannot learn the secret random angles $\{\theta_j\}_{j=1}^{N}$. One might naively think that this protocol can be generalized to the AKLT case by using the qutrit teleportation. However, again because of the non-trivial connection between the correlation space and physical particles, such a direct generalization does not work: because of the no-signalling principle a teleportation causes an error on the teleported state. Although such an error is not harmful in the qubit case [1], such an error in the qutrit teleportation destroys the particular structure of the correlation space required for correct computation [5]. Thus, a more elaborate two server scheme is needed. We now present a different type of the two-server scheme, where we can perform energy-gap-protected two server AKLT blind quantum computation with the AKLT Hamiltonian.

There are two key points which permit us to do this. Firstly, instead of teleporting a qutrit, Bob_1 teleports two qubits which are created from the application of the inverse PEPS $P^\dagger$ on the qutrit. The errors of the qubit teleportation do not destroy the required structure of the correlation space. Secondly, the extra rotation $Z_{\omega+\xi}$ by the angle $\omega + \xi$ which comes from the non-trivial

commutation relations between the PEPS and the teleportation errors, completely conceals the true computational angle ϕ from Bob_2.

The outline of our protocol is as follows. The single-qubit Z-rotation Z_ϕ in the correlation space is implemented in the following way: (I) Bob_1 creates (normal, not encrypted) AKLT states (Figs. 3.3 (a) and 3.4 (a)). (II) Bob_1 adiabatically turns off the interaction between a qutrit and the rest of qutrits in his resource state, and applies the inverse of PEPS $P^\dagger$ to the isolated qutrit (Figure 3.4 (b)) in order to convert the qutrit into a pair of two qubits. Bob_1 teleports these two qubits to Bob_2 by consuming Bell pairs (Figure 3.4 (c)). Because of the teleportation, the teleported qubits are affected by Pauli errors (Figure 3.4 (d)). Bob_1 also sends Alice the result of the measurement in the teleportation through the classical channel (Figure 3.3 (b)). (III) Bob_2 applies the filtering $\{|\eta_1\rangle\langle\eta_1|, I - |\eta_1\rangle\langle\eta_1|\}$ to the teleported two qubits. If the result is $|\eta_1\rangle\langle\eta_1|$, trivial Pauli operation is implemented in their computation. In this case, back to (II). If the result is $I - |\eta_1\rangle\langle\eta_1|$, Bob_2 further applies PEPS to convert the two qubits into a qutrit. (IV) Alice calculates the angle in which the qutrit should be measured by using her classical computer, and sends it to Bob_2 (Figure 3.3 (c)). (V) Bob_2 performs the measurement on the qutrit in that angle and sends the result of the measurement to Alice (Figure 3.3 (d)). Alice, Bob_1, and Bob_2 repeat (II)-(V) sufficiently many times. Now the operation $Z_{\omega+\xi+\phi}$ is implemented in the correlation space. As noted, the extra rotation $Z_{\omega+\xi}$ comes from the non-trivial commutation relations between PEPS P and the byproduct errors of the teleportation. Alice asks Bob_1 to compensate the byproduct operation $Z_{\omega+\xi}$ (Figure 3.3 (e) and (f)). Thus the desired Z-rotation Z_ϕ is finally implemented. The single-qubit X-rotation can be done in a similar way. We can show that two Bobs learn nothing about Alice's input, output, and algorithm. (For details see 3.6.3.)

This concludes the presentation of basic ideas of our scheme, and the rather technical details are left for a separate Section 3.6.3.

3.5 Discussion

In this chapter we have investigated the robustness of the basic ideas of the UBQC. In particular, we considered alternative implementations of Bob which may have advantages to the original graph-state based MBQC utilized in UBQC. We have focused on the adaptation of UBQC to AKLT state generalized measurement-based quantum computation In this computational model computation Bob performs could be protected from decoherence by an energy gap. While we have shown that it is possible to adapt UBQC to an AKLT-like setting, in the process we had to sacrifice the ground state properties of AKLT computation, which was exactly the initial motivation. The reason for this is simple – no large "encrypted state" can possibly be a ground state of a natural Hamiltonian, as the ground space we are interested in would have to contain all the possible encrypted states – *i.e.* exponentially many of them.

Thus, any single server blind computing scheme, which allows the server to maintain of his computation in a ground state of a natural Hamiltonian will have to be very different from UBQC.

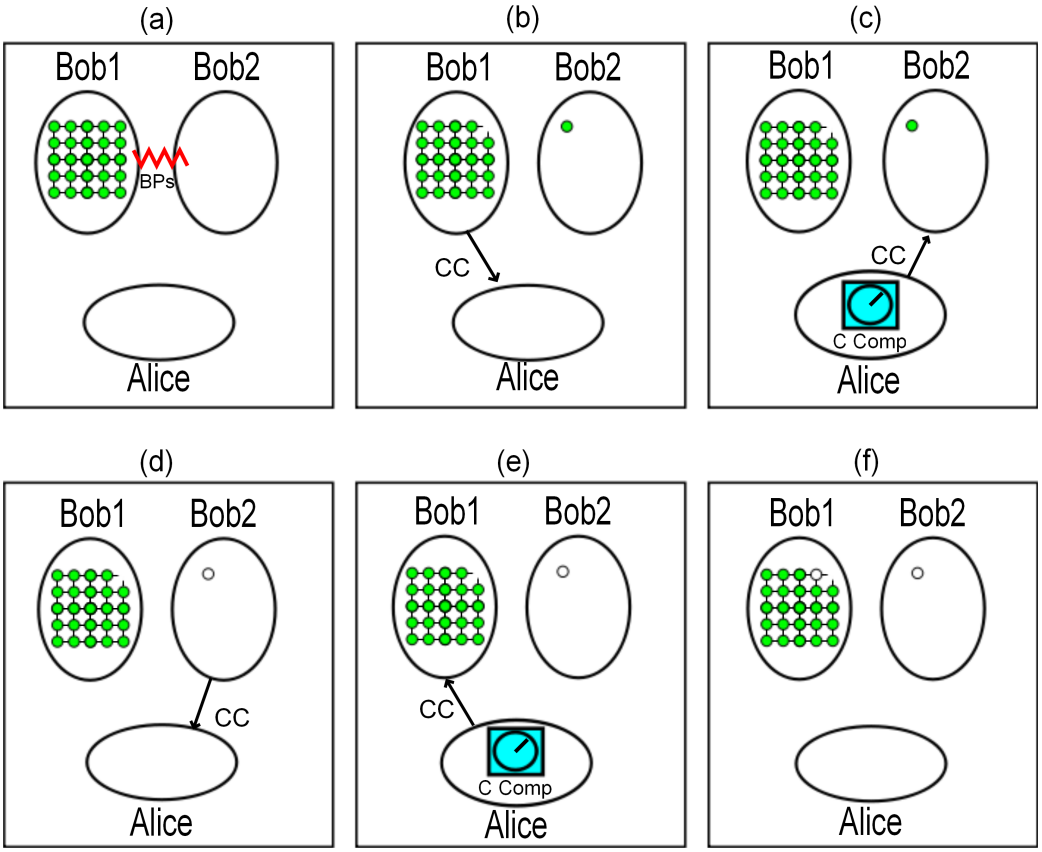

Figure 3.3. The two server protocol.

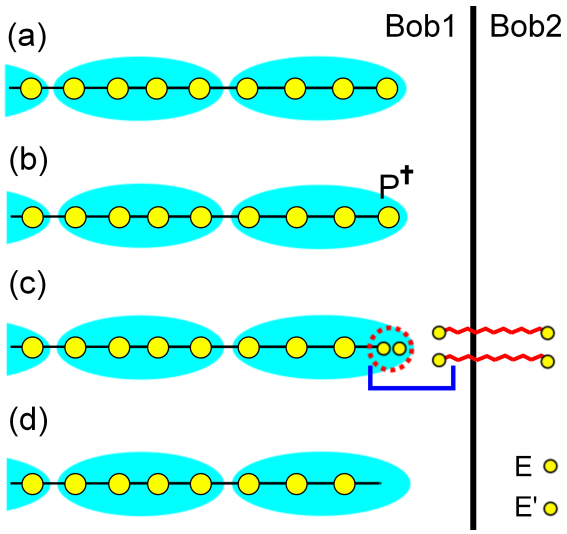

Figure 3.4. Teleportation from Bob_1 to Bob_2.

Interestingly, this issue disappears as soon as a two non-communicating server setting is introduced. In this setting, for the price of the no-communication assumption, both the ground state properties of the computation, and a fully classical client become possible. On one side, it is relatively clear why this works – in the two-server scheme we have presented, one of the servers fulfils the function of an extension of Alice herself, and does the measurements for her, while the other maintains the resource state in the ground state. The no-communication assumption thus protects Alice's privacy, making the measuring server effectively play the part of a trusted third party. However, this type of reasoning holds for the original two-server setting, presented in [1] as well. In their scheme, the first server is used to generate the encrypted state in the possession of the second. Interesting twist we have presented lays in the more active role the second server plays in our scheme – he does the actual measurements, and no encrypted resource is required. For this reason, the maintaining of the computation in the ground state was possible.

Similar "role inversion" can be performed in the single server setting as well. This idea inspired the "measurement-only Alice" UBQC variants [76], where the server only prepares the resource state (say, the cluster state) and sends qubit by qubit to Alice, who does measurements only, and no classical communication is present. In this setting also the resource Bob generates is not encrypted and thus can be in a ground state, and privacy is guaranteed by no-signalling. It may be possible to adapt the measurement only scheme to work with AKLT states (by having Alice assume the role of the second server), however in this case some communication between Alice and the server will have to exist, thus the security guarantees do not reduce trivially to no-signalling. This question we leave for future research.

3.5.1 Two-server setting and practice

Motivated by the double benefit of the two server setting, one may wonder can the non-communication assumption ever be justified in practice? In this section we present an attempt to create a somewhat realistic setting in which such an assumption could possibly be justified. We advise the reader not to take the following story overly seriously – it's main purpose is show that a no-communication assumption we use is not trivially and obviously never realizable in practice, and hopefully inspire a more studious approach to this problem.

In our imaginary setting many competing companies offer services of quantum computation. We consider $N + 1$ companies offering UBQC services and a trusted third party. The trusted third party acts as a control centre of the market, ensuring the companies behave honestly, and it alone has the power to eliminate a company from the market. We assume the control centre can prepare Bell pairs, and send them to the companies, and also that it has a private and authenticated channel to each of the companies and Alice. To begin the protocol, Alice contacts the control centre who then distributes a sufficient amount of Bell pairs to all the companies. Only the centre knows for each Bell pair which two companies share it. The centre then chooses randomly two companies to perform the two-server UBQC with, informing them individually over private channels. The centre keeps its choice secret. Then, the centre announces to all companies that a computation with Alice has begun. The two chosen companies we call players, and the other companies

non-players. The following two rules govern how a company is removed by the market:

1. If a non-player company A correctly guesses, and declares to the centre that another company B is a player, company B is eliminated from the market.

2. If a non-player company A incorrectly accuses a company B of being a player to the centre, *both* A and B are eliminated from the market.

We assume the companies all individually have the following goals, given in order of precedence: first, no company will behave in a way which leads to it being eliminated from the market. Second, every company has an interest in causing every other company to be eliminated from the market. Crucially, we assume the companies have more to gain by eliminating other companies from the market than by cheating on Alice. Under these assumptions no player company will try to cheat on Alice: in order to do so, they need to contact the other player company. But with probability $1 - 1/N$ they will choose a non-player company who then has the power to eliminate them from the market. Additionally, no non-player company will attempt to eliminate another company - they would have to correctly guess a player, since only they can be eliminated, and this will happen only with probability $2/N$. Note, if they guess incorrectly, they get eliminated too. If the assumptions on the agendas of companies are fulfilled, then no company will ever try to cheat on Alice, in fear of loosing its access to the market.

While it really is not a worthy challenge to find flaws in the scheme above, to the inspired reader we issue one which is: find a realistic setting which justifies the no-communication assumption, or prove this can never be done.

3.6 Technical details

3.6.1 Single-server blind quantum computing protocol

As said before, in the single-server BQC protocol (Figure 3.1), Alice has a classical computer and a quantum instrument that emits random four-qubit states which we will call "Dango states". Depending on the desired computation and the input size, Alice will send $(2 \times N \times M)$ Dango states directly to Bob through a one-way quantum channel that they initially share. Bob stores all of them in his quantum memory to creates the resource state, called "rotated AKLT states". The procedure of preparing such an initial state is explained in the Blind state preparation Subsection below.

Next, Alice calculates the angle in which a particle of a rotated AKLT state should be measured. Recall that this is a qutrit measurement that will induce a qubit operation over the correlation space. Moreover, the calculated angle should compensate for the initial random rotation of the Dango states and byproduct operation of the previous measurement. Finally, an additional random rotation will be added to hide the true result of the measurement from Bob. Bob performs the measurement according to Alice's information (sent via a classical channel to him), and returns the result of the measurement to Alice. They repeat this two-way classical communication until they finish the computation. Bob finally sends the final output of the quantum computation

to Alice. The exact protocol is given in the Blind computation Subsection below where we describe how a blind arbitrary X and Z rotation in the correlation space can be performed. Next we describe how two-qubit entangling operation of $\wedge Z$ can be performed in regular places. The rotation operators are also performed in regular interval, and hence the overall structure of the actual underlying computation remains hidden to Bob. These set of operators define a universal set of gates for quantum computing.

3.6.1.1 Preparation of the encrypted resource state

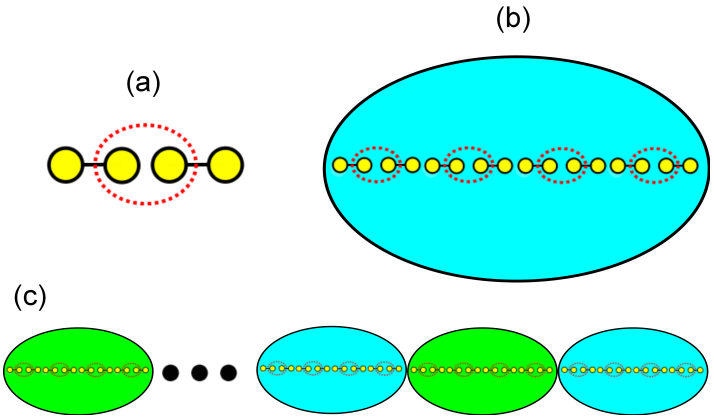

Figure 3.5. (a): A Dango state. Two yellow circles connected by a bond are qubits are in a Bell state. The operator $(I - |\eta_1\rangle\langle\eta_1|)T_{Z/X}(\theta)$ acts on two qubits in the red dotted circle. (b): A Z-Dango chain for $n = 4$. (c): A Combo chain $|C_b\rangle$. Z-Dango chains are colored in blue, whereas X-Dango chains are colored in green.

Denote the Bell basis with

$$\begin{aligned}
|\eta_1\rangle &\equiv \frac{1}{\sqrt{2}}\Big(|0\rangle\otimes|0\rangle + |1\rangle\otimes|1\rangle\Big),\\
|\eta_2\rangle &\equiv \frac{1}{\sqrt{2}}\Big(|0\rangle\otimes|0\rangle - |1\rangle\otimes|1\rangle\Big),\\
|\eta_3\rangle &\equiv \frac{1}{\sqrt{2}}\Big(|1\rangle\otimes|0\rangle + |0\rangle\otimes|1\rangle\Big),\\
|\eta_4\rangle &\equiv \frac{1}{\sqrt{2}}\Big(|1\rangle\otimes|0\rangle - |0\rangle\otimes|1\rangle\Big).
\end{aligned}$$

The full state preparation is described in Protocol 1. Assume Alice's desired computation is composed of a sequence of X and Z-rotations and $\wedge Z$ operations. Depending on the number of the required operators and the size of the input, Alice will choose integer values N and M. Then Alice's quantum instrument emits $N \times M$ "Z-Dango states" and $N \times M$ "X-Dango states" defined as (see Figure 3.5 (a))

$$\begin{aligned}
|D_Z(\theta^Z_{a,b})\rangle &\equiv (I\otimes(I-|\eta_1\rangle\langle\eta_1|)\otimes I)(I\otimes T_Z(\theta^Z_{a,b})\otimes I)|\eta_1\rangle\otimes|\eta_1\rangle,\\
|D_X(\theta^X_{a,b})\rangle &\equiv (I\otimes(I-|\eta_1\rangle\langle\eta_1|)\otimes I)(I\otimes T_X(\theta^X_{a,b})\otimes I)|\eta_1\rangle\otimes|\eta_1\rangle,
\end{aligned}$$

where $(a,b) \in \{1,...,N\} \times \{1,...,M\}$, $\theta_{a,b}^{Z/X} \in \mathcal{A} \equiv \left\{ \frac{k\pi}{4} \middle| k = 0,...,7 \right\}$ are independently and uniformly distributed random numbers which are secret to Bob, and the two qubit operators $T_Z(\theta_{a,b}^Z)$ and $T_X(\theta_{a,b}^X)$ are defined by

$$\begin{aligned} T_Z(\theta_{a,b}^Z) &\equiv |00\rangle\langle 00| + e^{i\theta_{a,b}^Z}|01\rangle\langle 01| + |10\rangle\langle 10| + |11\rangle\langle 11|, \\ T_X(\theta_{a,b}^X) &\equiv \Big(\frac{1+e^{i\theta_{a,b}^X}}{2}|\eta_2\rangle + \frac{1-e^{i\theta_{a,b}^X}}{2}|\eta_4\rangle\Big)\langle\eta_2| + \\ &\quad \Big(\frac{1-e^{i\theta_{a,b}^X}}{2}|\eta_2\rangle + \frac{1+e^{i\theta_{a,b}^X}}{2}|\eta_4\rangle\Big)\langle\eta_4| + |\eta_1\rangle\langle\eta_1| + |\eta_3\rangle\langle\eta_3|. \end{aligned}$$

Alice sends all these Dango states to Bob, and records all $\{\theta_{a,b}^{Z/X}\}$ for the later use. Bob arranges all the Dango states in a lattice with $2N$ columns and M rows.

Alice chooses a parameter $n < N$. We call a collection of n Dango states, sent by Alice to be kept in Bob's memory, " (k,b)th Z-Dango chain states" or "(k,b)th X-Dango chain states" defined as (see Figure 3.5 (b))

$$\begin{aligned} |B_{k,b}^Z\rangle &\equiv \bigotimes_{j=1}^{n} |D_Z(\theta_{(k-1)n+j,b}^Z)\rangle, \\ |B_{k,b}^X\rangle &\equiv \bigotimes_{j=1}^{n} |D_X(\theta_{(k-1)n+j,b}^X)\rangle, \end{aligned}$$

where $k = 1,...,N/n$ and $b = 1,...,M$. A Z-Dango chain state is used for the implementation of a single-qubit Z-rotation whereas an X-Dango chain state is used for the implementation a single-qubit X-rotation. However, to hide the actual structure of the computation Alice will work with a regular one-dimensional chain, called "Combo chain state" $|C_b\rangle$, composed of N/n Z-Dango chain states and N/n X-Dango chain states (Figure 3.5 (c)) with two-edge qubits projected on $|R^*\rangle$ and $|L\rangle$ states, respectively (Figure 3.2 (a)):

$$|C_b\rangle \equiv \langle R^*|\langle L|\Big(|B_{N/n,b}^X\rangle \otimes |B_{N/n,b}^Z\rangle \otimes ... \otimes |B_{2,b}^X\rangle \otimes |B_{2,b}^Z\rangle \otimes |B_{1,b}^X\rangle \otimes |B_{1,b}^Z\rangle\Big)$$

($b = 1,...,M$). Here, $|L\rangle = (|0\rangle + i|1\rangle)/\sqrt{2}$ and $|R\rangle = |0\rangle$.

However, $|R\rangle$ could be any arbitrary state depending on Alice's desired input using the an uploading method [77] [4].

Now in order to entangle qubits of the Combo chain to create the desired resource state, Bob will perform the following operators. Recall, the PEPS operation is defined with [66, 71]

$$P \equiv \frac{1}{\sqrt{2}} \sum_{l=1}^{3} \sum_{i=0}^{1} \sum_{j=0}^{1} A_{i,j}[l]|l\rangle\langle i| \otimes \langle j|$$

and creates a qutrit from two qubits, where $A_{i,j}[l]$ is (i,j)-element of the matrix $A[l]$. Consider

[4] It is possible to "upload" a quantum state into the correlation space by using a teleportation-like procedure. If the state has been properly quantum one-time padded before the uploading procedure, unconditionally blind computation can be performed on this input using similar methods for quantum input presented in [1].

the following unitary operators acting on a qutrit

$$\begin{aligned} U(\theta_{a,b}^{Z/X}) &\equiv \Big(\frac{1+e^{i\theta_{a,b}^{Z/X}}}{2}|1\rangle + \frac{1-e^{i\theta_{a,b}^{Z/X}}}{2}|2\rangle\Big)\langle 1| + \\ &\quad \Big(\frac{1-e^{i\theta_{a,b}^{Z/X}}}{2}|1\rangle + \frac{1+e^{i\theta_{a,b}^{Z/X}}}{2}|2\rangle\Big)\langle 2| + |3\rangle\langle 3|, \\ V &\equiv |3\rangle\langle 1| + |1\rangle\langle 2| + |2\rangle\langle 3|. \end{aligned}$$

It is easy to verify that

$$\begin{aligned} PT_Z(\theta_{a,b}^{Z/X}) &= U(\theta_{a,b}^{Z/X})P, \\ PT_X(\theta_{a,b}^{Z/X}) &= V^\dagger U(\theta_{a,b}^{Z/X})VP. \end{aligned}$$

Bob has to apply the filtering operation $I - |\eta_1\rangle\langle\eta_1|$. In order to do so, he performs the measurement $\{|\eta_1\rangle\langle\eta_1|, I - |\eta_1\rangle\langle\eta_1|\}$ to every pair of two qubits in the Combo chain which is specified by a dotted blue circle in Figure 3.2 (b). If $|\eta_1\rangle\langle\eta_1|$ is realised, two qubits are just removed from the chain (Figure 3.2 (c)). Next Bob applies the PEPS operation P to each pair of two qubits in order to obtain qutrits (Figure 3.2 (d), (e), (f)). This PEPS operation is done deterministically since $I - |\eta_1\rangle\langle\eta_1|$ is already applied to every pair of qubits. Therefore, Bob has created a new one-dimensional chain of qutrits (Figure 3.2 (f)) called "rotated AKLT state":

$$|RAKLT_b^{2N,L,R}(\{\theta_{a,b}^{Z/X}\})\rangle = \mathcal{U}_b(\{\theta_{a,b}^{Z/X}\})|AKLT^{2N,L,R}\rangle,$$

where

$$\begin{aligned} \mathcal{U}_b(\{\theta_{a,b}^{Z/X}\}) \equiv & \Big\{V^\dagger U(\theta_{N,b}^X)V \otimes ... \otimes V^\dagger U(\theta_{N-n+1,b}^X)V\Big\} \otimes \\ & \Big\{U(\theta_{N,b}^Z) \otimes ... \otimes U(\theta_{N-n+1,b}^Z)\Big\} \\ & \vdots \\ & \otimes\Big\{V^\dagger U(\theta_{n+n,b}^X)V \otimes ... \otimes V^\dagger U(\theta_{n+1,b}^X)V\Big\} \otimes \\ & \Big\{U(\theta_{n+n,b}^Z) \otimes ... \otimes U(\theta_{n+1,b}^Z)\Big\} \otimes \\ & \Big\{V^\dagger U(\theta_{n,b}^X)V \otimes ... \otimes V^\dagger U(\theta_{1,b}^X)V\Big\} \otimes \Big\{U(\theta_{n,b}^Z) \otimes ... \otimes U(\theta_{1,b}^Z)\Big\} \end{aligned}$$

(for simplicity, we have assumed that all filterings give $|\eta_1\rangle\langle\eta_1|$). We call a qutrit which is rotated by $U(\theta_{a,b}^Z)$ "Z-prerotated qutrit" and a qutrit which is rotated by $V^\dagger U(\theta_{a,b}^X)V$ "X-prerotated qutrit". Other qutrits are called "plain qutrits". The "(k,b)th Z/X-prerotated AKLT subsystem" is defined to be the set of Z/X-prerotated qutrits in bth prerotated AKLT chain corresponding to particles of (k,b)th Z/X-Dango chain.

Next we show how the actual blind computation is performed, and prove the security of the scheme.

Protocol 7 The preparation of the encrypted resource state

- Alice sends Bob parameter values N, M and $n < N$.
- Alice sends Bob $N \times M$ many Z-Dango states $|D_Z(\theta^Z_{a,b})\rangle$ where $(a, b) \in \{1, ..., N\} \times \{1, ..., M\}$.
- Alice sends Bob $N \times M$ many X-Dango states $|D_X(\theta^X_{a,b})\rangle$ where $(a, b) \in \{1, ..., N\} \times \{1, ..., M\}$.
- Bob arranges the received Z-Dango states in M rows of N/n Z-Dango chains $|B^Z_{k,b}\rangle$ where $k = 1, 2, ..., N/n$, $b = 1, ..., M$.
- Bob arranges the received X-Dango states in M rows of N/n X-Dango chains $|B^X_{k,b}\rangle$ where $k = 1, 2, ..., N/n$, $b = 1, ..., M$.
- Bob arranges the Dango chains in M rows of Combo chains $|C_b\rangle \equiv |B^X_{N/n,b}\rangle \otimes |B^Z_{N/n,b}\rangle \otimes ... \otimes |B^X_{1,b}\rangle \otimes |B^Z_{1,b}\rangle$.
- Bob applies filtering and PEPS operators to create M rows of rotated AKLT states $|RAKLT^{2N,L,R}_b(\{\theta^{Z/X}_{a,b}\})\rangle$, where $b = 1, ..., M$.

See also Figure 3.2.

3.6.1.2 Single-server blind quantum computation protocol

A single blind Z-rotation is performed using a Z-Dango chain state. Assume that Alice wants to perform the Z-rotation $\exp\left[\frac{iZ}{2}\phi^Z_{k,b}\right]$ with $\phi^Z_{k,b} \in \mathcal{A}$, using the (k, b)th Z-Dango chain.

Recall, $\exp\left[\frac{iZ}{2}\phi\right] = Z_\phi$, up to a global phase, and in this section we maintain the first notation to avoid excessive nested indices. The steps of this process are given in Protocol 2. Note that we implement the desired Z-rotation using qutrit measurements. However, the third outcome $|\gamma\rangle$ leads to the failure as it implements only the trivial Pauli Z. The probability that Alice fails to implement her desired Z-rotation in a single Z-Dango chain is $1/3^n$, which can be made arbitrarily small by choosing a sufficiently large n. Similarly a blind X rotation could be applied where again, the probability that Alice fails to implement her desired X-rotation in a single X-Dango chain is $1/3^n$. Note that these protocols are designed in such a way to follow the construction of the AKLT computation, while canceling the prerotation that was added to the resource state for the purpose of blindness.

Next, we explain how the two-qubit operation of controlled-Z ($\wedge Z$) is performed. In order to perform $\wedge Z$ gates blindly, $\wedge Z$ gates are periodically implemented with the period that is independent of Alice's input and the algorithm. In this case, Bob learns nothing from the period. Because of the periodic implementation of $\wedge Z$ gates, Alice sometimes experiences an unwanted $\wedge Z$ gate. However, Alice can cancel the effect of an unwanted $\wedge Z$ gate by implementing the trivial identity operation (plus Pauli byproduct operations) until she arrives at the next $\wedge Z$ gate which cancels the previous one. The commutation rules which allow for this are the following:

$$\begin{aligned}
\wedge Z(I \otimes X) \wedge Z &= Z \otimes X, \\
\wedge Z(I \otimes Z) \wedge Z &= I \otimes Z, \\
\wedge Z(X \otimes I) \wedge Z &= X \otimes Z, \\
\wedge Z(Z \otimes I) \wedge Z &= Z \otimes I.
\end{aligned}$$

Protocol 8 Blind Z Rotation

Initially the flag parameter (known to both Alice and Bob) is set $\tau = 1$. For $j = 1 \cdots n$ Alice and Bob perform the following steps.

(I) Alice sends Bob the angle

$$\delta^Z_{(k-1)n+j,b} \equiv \tau\phi^Z_{k,b} + \theta^Z_{(k-1)n+j,b} + \pi r^Z_{(k-1)n+j,b} \pmod{2\pi},$$

where $r^Z_{(k-1)n+j,b} \in \{0,1\}$ is chosen uniformly at random and unknown to Bob. If there is the X byproduct before this step, $\phi^Z_{k,b}$ should be replaced with $-\phi^Z_{k,b}$ in order to compensate this byproduct operator. However, Z byproduct commutes trivially with the operation implemented in the correlation space, and therefore it can be corrected at the end of computation.

(II) Bob measures the jth Z-prerotated qutrit of the (k,b)th Z-prerotated AKLT subsystem in the basis

$$\mathcal{M}(\delta^Z_{(k-1)n+j,b}) \equiv \{|\alpha(\delta^Z_{(k-1)n+j,b})\rangle, |\beta(\delta^Z_{(k-1)n+j,b})\rangle, |\gamma\rangle\},$$

where

$$\begin{aligned}
|\alpha(\delta^Z_{(k-1)n+j,b})\rangle &\equiv \frac{1+\exp\left[i\delta^Z_{(k-1)n+j,b}\right]}{2}|1\rangle + \frac{1-\exp\left[i\delta^Z_{(k-1)n+j,b}\right]}{2}|2\rangle, \\
|\beta(\delta^Z_{(k-1)n+j,b})\rangle &\equiv \frac{1-\exp\left[i\delta^Z_{(k-1)n+j,b}\right]}{2}|1\rangle + \frac{1+\exp\left[i\delta^Z_{(k-1)n+j,b}\right]}{2}|2\rangle, \\
|\gamma\rangle &\equiv |3\rangle.
\end{aligned}$$

and sends the result to Alice.

- If the measurement result is $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$,

$$R^\alpha_Z(\tau\phi^Z_{k,b}, r^Z_{(k-1)n+j,b}) \equiv \exp\left[\frac{-i\tau\phi^Z_{k,b}}{2}\right] X Z^{r^Z_{(k-1)n+j,b}} \exp\left[\frac{iZ}{2}\tau\phi^Z_{k,b}\right]$$

  is implemented in the correlation space and Alice sets $\tau = 0$.
- If the measurement result is $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$,

$$R^\beta_Z(\tau\phi^Z_{k,b}, r^Z_{(k-1)n+j,b}) \equiv \exp\left[\frac{-i\tau\phi^Z_{k,b}}{2}\right] X Z^{r^Z_{(k-1)n+j,b}+1} \exp\left[\frac{iZ}{2}\tau\phi^Z_{k,b}\right]$$

  is implemented in the correlation space and Alice sets $\tau = 0$.
- If the measurement result is $|\gamma\rangle$, Z is implemented in the correlation space.

The probability of obtaining each result is $1/3$.

(III) Bob does the measurement $\{|1\rangle, |2\rangle, |3\rangle\}$ on the next plain (not pre-rotated) qutrit if any.

Protocol 9 3: Blind X Rotation

Initially the flag parameter (known to both Alice and Bob) is set $\tau = 1$. For $j = 1 \cdots n$ Alice and Bob perform the following steps.

(I) Alice sends Bob the angle

$$\delta^X_{(k-1)n+j,b} \equiv \tau\phi^X_{k,b} + \theta^X_{(k-1)n+j,b} + r^X_{(k-1)n+j,b}\pi \pmod{2\pi},$$

where $r^X_{(k-1)n+j,b} \in \{0,1\}$ is chosen uniformly at random (and unknown to Bob). If a Z byproduct operator occurred before this step, $\phi^X_{k,b}$ should be replaced with $-\phi^X_{k,b}$ in order to compensate this byproduct operator. However, X byproduct commutes trivially with the operation implemented in the correlation space, and therefore it can be corrected at the end of computation.

(II) Bob applies V on the jth X-prerotated qutrit of the (k,b)th X-prerotated AKLT subsystem, and does the measurement in the basis

$$\mathcal{M}(\delta^X_{(k-1)n+j,b}) \equiv \{|\alpha(\delta^X_{(k-1)n+j,b})\rangle, |\beta(\delta^X_{(k-1)n+j,b})\rangle, |\gamma\rangle\},$$

where

$$\begin{aligned}
|\alpha(\delta^X_{(k-1)n+j,b})\rangle &\equiv \frac{1+\exp\left[i\delta^X_{(k-1)n+j,b}\right]}{2}|1\rangle + \frac{1-\exp\left[i\delta^X_{(k-1)n+j,b}\right]}{2}|2\rangle,\\
|\beta(\delta^X_{(k-1)n+j,b})\rangle &\equiv \frac{1-\exp\left[i\delta^X_{(k-1)n+j,b}\right]}{2}|1\rangle + \frac{1+\exp\left[i\delta^X_{(k-1)n+j,b}\right]}{2}|2\rangle,\\
|\gamma\rangle &\equiv |3\rangle.
\end{aligned}$$

and sends the result to Alice.

- If the measurement result is $|\alpha(\delta^X_{(k-1)n+j,b})\rangle$,

$$R^\alpha_X(\tau\phi^X_{k,b}, r^X_{(k-1)n+j,b}) \equiv \exp\left[\frac{-i\tau\phi^X_{k,b}}{2}\right] X^{r^X_{(k-1)n+j,b}+1} Z \exp\left[\frac{-iX}{2}\tau\phi^X_{k,b}\right]$$

  is implemented in the correlation space and Alice sets $\tau = 0$.
- If the measurement result is $|\beta(\delta^X_{(k-1)n+j,b})\rangle$,

$$R^\beta_X(\tau\phi^X_{k,b}, r^X_{(k-1)n+j,b}) \equiv \exp\left[\frac{-i\tau\phi^X_{k,b}}{2}\right] X^{r^X_{(k-1)n+j,b}} Z \exp\left[\frac{-iX}{2}\tau\phi^X_{k,b}\right]$$

  is implemented in the correlation space and Alice sets $\tau = 0$.
- If the measurement result is $|\gamma\rangle$, X is implemented.

The probability of obtaining each result is $1/3$.

(III) Bob performs the measurement $\{|1\rangle, |2\rangle, |3\rangle\}$ on the next plain qutrit if any.

In Protocol 4, we explain how to implement the $\wedge Z$ gate plus Z-rotation

$$\left(\exp\left[\frac{iZ}{2}\phi^Z_{k,b}\right] \otimes \exp\left[\frac{iZ}{2}\phi^Z_{k',b'}\right]\right) \wedge Z$$

between (k,b)th and (k',b')th Z-prerotated AKLT subsystems, where $\phi^Z_{k,b}, \phi^Z_{k',b'} \in \mathcal{A}$. Note that the local Z rotations are required to cancel the prerotations of qutrits. The probability that Alice fails to implement the $\wedge Z$ gate in this algorithm is $(5/9)^n$, which is arbitrarily small for n chosen sufficiently large.

3.6.2 Proof of blindness of the single-server protocol

In this section, we show the blindness of the single-server protocol (composed of Protocols 8, 9, 10). As we have already seen, informally speaking, a protocol is defined to be blind if Bob, given all the classical and quantum information during the protocol, cannot learn anything about the Alice's actual computational angles, input and the output [1]. In the original paper for the blind quantum computation over the cluster states [1] blindness is formally defined in terms of the independence of classical and quantum states Bob receives throughout the protocol from Alice's secret. Here, we adapt the definition to our setting but the two definitions can be shown to be equivalent [5].

Definition 1. *A single-server protocol is blind if*

(S1) The conditional probability distribution of Alice's nontrivial computational angles, given all the classical information Bob can obtain during the protocol, and given the measurement results of any POVMs which Bob may perform on his system at any stage of the protocol, is uniform,

and

(S2) The register state in the correlation space is quantum one-time padded.

In order to show (S1), we have to show three lemmas.

In the following we define Δ, Φ, Θ and R to be independently and uniformly distributed random variables, corresponding to the angles sent by Alice to Bob, Alice's secret computational angle, random prerotation and a hidden binary parameter, respectively. From the construction of the protocol, the following relation is satisfied:

$$\Delta = \Phi + \Theta + R\pi \pmod{2\pi}.$$

We denote by ρ_Θ the state that Alice sends to Bob parametrized by Θ. The most general method Bob may resort to in order to learn Alice's secret computational angles is described by a POVM

[5] In the original definition, it is required that the state of Bob's register does not depend on the computational angles. Here, we will show that the probability distribution describing Bob's state of knowledge about the computational angles remains the same even when conditioned on any outcome of any POVM Bob may apply on his system. This means that, for every POVM, the measurement outcomes are independent from Alice's angles. But this means the states of Bob register must be independent from the angles.

Protocol 10 Controlled-Z followed by Blind Z-rotations

Initially the flag parameters (known to both Alice and Bob) are set $\tau = 1$, $\tau' = 1$ and $\omega = 1$. For $j = 1 \cdots n$ Alice and Bob perform the following steps.

(I) If $\omega = 0$, skip this step. Bob applies the unitary operation

$$\begin{aligned} W \equiv & \frac{|1,1\rangle + |1,2\rangle + |2,1\rangle - |2,2\rangle}{2}\langle 1,1| \\ & + \frac{|1,1\rangle + |1,2\rangle - |2,1\rangle + |2,2\rangle}{2}\langle 1,2| \\ & + \frac{|1,1\rangle - |1,2\rangle + |2,1\rangle + |2,2\rangle}{2}\langle 2,1| \\ & + \frac{-|1,1\rangle + |1,2\rangle + |2,1\rangle + |2,2\rangle}{2}\langle 2,2| \\ & + |1,3\rangle\langle 1,3| + |2,3\rangle\langle 2,3| + |3,1\rangle\langle 3,1| + |3,2\rangle\langle 3,2| + |3,3\rangle\langle 3,3| \end{aligned}$$

between jth Z-prerotated qutrit of (k,b)th Z-prerotated AKLT subsystem and jth Z-prerotated qutrit of (k',b')th Z-prerotated AKLT subsystem.

(II) Alice sends Bob the angles

$$\begin{aligned} \delta^Z_{(k-1)n+j,b} &= \tau\phi^Z_{k,b} + \theta^Z_{(k-1)n+j,b} + r^Z_{(k-1)n+j,b}\pi \pmod{2\pi}, \\ \delta^Z_{(k'-1)n+j,b'} &= \tau'\phi^Z_{k',b'} + \theta^Z_{(k'-1)n+j,b'} + r^Z_{(k'-1)n+j,b'}\pi \pmod{2\pi}, \end{aligned}$$

where $r^Z_{(k-1)n+j,b}, r^Z_{(k'-1)n+j,b'} \in \{0,1\}$ are random numbers. If there is any byproduct which contains X before this step, the sign of $\phi^Z_{k,b}$ or $\phi^Z_{k',b'}$ should be appropriately changed. However, Z byproduct commutes trivially with the operation implemented in the correlation space, and therefore it can be corrected at the end of computation.

(III) Bob does the measurement $\mathcal{M}(\delta^Z_{(k-1)n+j,b})$ (the same as that of Protocol 1) on the jth Z-prerotated qutrit of the (k,b)th Z-prerotated AKLT subsystem and the measurement $\mathcal{M}(\delta^Z_{(k'-1)n+j,b'})$ on the jth Z-prerotated qutrit of the (k',b')th Z-prerotated AKLT subsystem. The operation implemented in the correlation space is summarized as follows:

- $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle \otimes |\alpha(\delta^Z_{(k'-1)n+j,b'})\rangle$: $\left[R^\alpha_Z(\tau\phi^Z_{k,b}, r^Z_{(k-1)n+j,b}) \otimes R^\alpha_Z(\tau'\phi^Z_{k',b'}, r^Z_{(k'-1)n+j,b'})\right] \wedge Z^\omega$
- $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle \otimes |\beta(\delta^Z_{(k'-1)n+j,b'})\rangle$: $\left[R^\alpha_Z(\tau\phi^Z_{k,b}, r^Z_{(k-1)n+j,b}) \otimes R^\beta_Z(\tau'\phi^Z_{k',b'}, r^Z_{(k'-1)n+j,b'})\right] \wedge Z^\omega$
- $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle \otimes |\gamma\rangle$: $R^\alpha_Z(\tau\phi^Z_{k,b}, r^Z_{(k-1)n+j,b}) \otimes Z$
- $|\beta(\delta^Z_{(k-1)n+j,b})\rangle \otimes |\alpha(\delta^Z_{(k'-1)n+j,b'})\rangle$: $\left[R^\beta_Z(\tau\phi^Z_{k,b}, r^Z_{(k-1)n+j,b}) \otimes R^\alpha_Z(\tau'\phi^Z_{k',b'}, r^Z_{(k'-1)n+j,b'})\right] \wedge Z^\omega$
- $|\beta(\delta^Z_{(k-1)n+j,b})\rangle \otimes |\beta(\delta^Z_{(k'-1)n+j,b'})\rangle$: $\left[R^\beta_Z(\tau\phi^Z_{k,b}, r^Z_{(k-1)n+j,b}) \otimes R^\beta_Z(\tau'\phi^Z_{k',b'}, r^Z_{(k'-1)n+j,b'})\right] \wedge Z^\omega$
- $|\beta(\delta^Z_{(k-1)n+j,b})\rangle \otimes |\gamma\rangle$: $R^\beta_Z(\tau\phi^Z_{k,b}, r^Z_{(k-1)n+j,b}) \otimes Z$
- $|\gamma\rangle \otimes |\alpha(\delta^Z_{(k'-1)n+j,b'})\rangle$: $Z \otimes R^\alpha_Z(\tau'\phi^Z_{k',b'}, r^Z_{(k'-1)n+j,b'})$
- $|\gamma\rangle \otimes |\beta(\delta^Z_{(k'-1)n+j,b'})\rangle$: $Z \otimes R^\beta_Z(\tau'\phi^Z_{k',b'}, r^Z_{(k'-1)n+j,b'})$
- $|\gamma\rangle \otimes |\gamma\rangle$: $Z \otimes Z$

(IV) If the Z-rotation by $\phi^Z_{k,b}$ is implemented in the preivous step, then Alice sets $\tau = 0$. If the z-rotation by $\phi^Z_{k',b'}$ is implemented in the previous step, then Alice sets $\tau' = 0$. If the $\wedge Z$ is implemented in the previous step, then Alice sets $\omega = 0$.

(V) Bob does the measurement $\{|1\rangle, |2\rangle, |3\rangle\}$ on the next plain qutrit if any.

measurement $\{\Pi_j\}_{j=1}^m$ on ρ_Θ. This POVM can depend on all classical messages received from Alice. Let $O \in \{1, ..., m\}$ be the random variable corresponding to the result of the POVM measurement. Bob's knowledge about Alice's secret angles is given by the conditional probability distribution of $\Phi = \phi$ given $O = j$ and $\Delta = \delta$:

$$P(\Phi = \phi | O = j, \Delta = \delta).$$

Lemma 1. *If ρ_Θ is a Dango state $|D_{Z/X}(\Theta)\rangle\langle D_{Z/X}(\Theta)|$, then $P(\Phi = \phi|O = j, \Delta = \delta) = \frac{1}{8}$ for any $\phi, \delta \in \mathcal{A}$, $j \in \{1, ..., m\}$, and POVM on ρ_Θ.*

Proof: From Bayes' theorem, we have

$$\begin{aligned} P(\Phi = \phi|O = j, \Delta = \delta) &= \frac{P(O = j|\Phi = \phi, \Delta = \delta)P(\Phi = \phi, \Delta = \delta)}{P(O = j, \Delta = \delta)} \\ &= \frac{P(O = j|\Phi = \phi, \Delta = \delta)P(\Phi = \phi)P(\Delta = \delta)}{P(O = j|\Delta = \delta)P(\Delta = \delta)} \\ &= \frac{1}{8}\frac{\mathrm{Tr}\big[\Pi_j \frac{1}{2}\sum_r \rho_{\delta-\phi-r\pi}\big]}{\mathrm{Tr}\big[\Pi_j \frac{1}{8}\frac{1}{2}\sum_{\phi,r} \rho_{\delta-\phi-r\pi}\big]}. \end{aligned}$$

If ρ_Θ is a Dango state $|D_{Z/X}(\Theta)\rangle\langle D_{Z/X}(\Theta)|$, we obtain

$$\frac{1}{2}\sum_r \rho_{\delta-\phi-r\pi} = \frac{1}{2}\frac{1}{8}\sum_{\phi,r} \rho_{\delta-\phi-r\pi}$$

for any $\delta, \phi \in [0, 2\pi]$, and hence $P(\Phi = \phi|O = j, \Delta = \delta) = 1/8$. The above equation is valid since

$$|D_Z(\theta)\rangle\langle D_Z(\theta)| =$$

$$\big[I \otimes (I - |\eta_1\rangle\langle\eta_1|) \otimes I\big]\frac{1}{4}\begin{pmatrix} 1 & e^{-i\theta} & 1 & 1 \\ e^{i\theta} & 1 & e^{i\theta} & e^{i\theta} \\ 1 & e^{-i\theta} & 1 & 1 \\ 1 & e^{-i\theta} & 1 & 1 \end{pmatrix} \oplus \mathbf{0}_{12}\big[I \otimes (I - |\eta_1\rangle\langle\eta_1|) \otimes I\big],$$

where $\mathbf{0}_{12}$ is the 12×12 zero matrix and the 4×4 matrix is in the basis $\{|0000\rangle, |0011\rangle, |1100\rangle, |1111\rangle\}$, and there exists a unitary which maps $(I \otimes T_Z(\theta) \otimes I)|\eta_1\rangle \otimes |\eta_1\rangle$ to $(I \otimes T_X(\theta) \otimes I)|\eta_1\rangle \otimes |\eta_1\rangle$. ■

Lemma 2. *Consider a collection of L states $\{\rho_{\Theta_l}\}_l$ such that for each $l = 1, \ldots, L$ we have*

$$P\Big(\Phi_l = \phi_l \Big| O^l = j, \Delta_l = \delta_l\Big) = \frac{1}{8}$$

where Δ_l, Φ_l, Θ_l and R_l are defined as before. Also, O^l are the random variables corresponding to (arbitrary) POVMs performed on individual systems. Then for any global POVM performed on the entire collection of L states we have

$$P\Big(\Phi = (\phi_1, ..., \phi_L)\Big|O = j, \Delta = (\delta_1, ..., \delta_L)\Big) = \frac{1}{8^L}$$

where O is the random variable corresponding to the outcome of the global POVM.

Proof: Similar to the previous proof, from Bayes' theorem, we have

$$\begin{aligned}
& P\Big(\Phi = (\phi_1, ..., \phi_L)\Big|O = j, \Delta = (\delta_1, ..., \delta_L)\Big) \\
= & \frac{P\Big(O = j\Big|\Phi = (\phi_1, ..., \phi_L), \Delta = (\delta_1, ..., \delta_L)\Big)P\Big(\Phi = (\phi_1, ..., \phi_L), \Delta = (\delta_1, ..., \delta_L)\Big)}{P\Big(O = j, \Delta = (\delta_1, ..., \delta_L)\Big)} \\
= & \frac{P\Big(O = j\Big|\Phi = (\phi_1, ..., \phi_L), \Delta = (\delta_1, ..., \delta_L)\Big)P\Big(\Phi = (\phi_1, ..., \phi_L)\Big)\Big)}{P\Big(O = j\Big|\Delta = (\delta_1, ..., \delta_L)\Big)} \times \\
& \frac{P\Big(\Delta = (\delta_1, ..., \delta_L)\Big)}{P\Big(\Delta = (\delta_1, ..., \delta_L)\Big)} \\
= & \frac{1}{8^L}\frac{\mathrm{Tr}\Big[\Pi_j \bigotimes_{i=1}^{L} \frac{1}{2}\sum_{r_i} \rho_{\delta_i - \phi_i - r_i\pi}\Big]}{\mathrm{Tr}\Big[\Pi_j \bigotimes_{i=1}^{L} \frac{1}{8}\frac{1}{2}\sum_{\phi_i, r_i} \rho_{\delta_i - \phi_i - r_i\pi}\Big]}.
\end{aligned}$$

Let us define two local operators acting on lth state ρ_{Θ_l} by

$$\begin{aligned}
\Pi_j^l &\equiv \mathrm{Tr}_{1,...,l-1,l+1,...,L}\Big[\Pi_j \bigotimes_{i=1}^{l-1}\frac{1}{2}(\sum_{r_i}\rho_{\delta_i-\phi_i-r_i\pi}) \bigotimes_{i=l+1}^{L}\frac{1}{2}(\sum_{r_i}\rho_{\delta_i-\phi_i-r_i\pi})\Big], \\
\tilde{\Pi}_j^l &\equiv \mathrm{Tr}_{1,...,l-1,l+1,...,L}\Big[\Pi_j \bigotimes_{i=1}^{l-1}\frac{1}{8}\frac{1}{2}(\sum_{\phi_i,r_i}\rho_{\delta_i-\phi_i-r_i\pi}) \bigotimes_{i=l+1}^{L}\frac{1}{8}\frac{1}{2}(\sum_{\phi_i,r_i}\rho_{\delta_i-\phi_i-r_i\pi})\Big].
\end{aligned}$$

The partial trace is a CPTP superoperator, hence the above operators are non-negative operators, and since

$$\begin{aligned}
\sum_{j=1}^{m}\Pi_j^l &= I, \\
\text{and also } \sum_{j=1}^{m}\tilde{\Pi}_j^l &= I,
\end{aligned}$$

it follows that $\{\Pi_j^l\}_{j=1}^m$ and $\{\tilde{\Pi}_j^l\}_{j=1}^m$ are local POVMs on lth state.

Let O^l and $\tilde{O}^l$ be the random variables which correspond to the results of the POVMs $\{\Pi_j^l\}_{j=1}^m$

and $\{\tilde{\Pi}^l_j\}_{j=1}^m$, respectively. Then, we have

$$\begin{aligned}
P\Big(\Phi=(\phi_1,...,\phi_L)\Big|O=j,\Delta=(\delta_1,...,\delta_L)\Big) = \\
\frac{1}{8^L}\frac{\mathrm{Tr}_l\Big[\Pi^l_j\frac{1}{2}\sum_{r_l}\rho_{\delta_l-\phi_l-r_l\pi}\Big]}{\mathrm{Tr}_l\Big[\tilde{\Pi}^l_j\frac{1}{8}\frac{1}{2}\sum_{\phi_l,r_l}\rho_{\delta_l-\phi_l-r_l\pi}\Big]} = \\
\frac{1}{8^L}\frac{P(O^l=j|\Delta_l=\delta_l,\Phi_l=\phi_l)}{P(\tilde{O}^l=j|\Delta_l=\delta_l)} = \\
\frac{1}{8^L}\frac{P(O^l=j|\Delta_l=\delta_l,\Phi_l=\phi_l)}{P(O^l=j|\Delta_l=\delta_l)}\frac{P(O^l=j|\Delta_l=\delta_l)}{P(\tilde{O}^l=j|\Delta_l=\delta_l)} = \\
\frac{1}{8^{L-1}}\frac{P(O^l=j|\Delta_l=\delta_l,\Phi_l=\phi_l)P(\Delta_l=\delta_l)P(\Phi_l=\phi_l)}{P(O^l=j|\Delta_l=\delta_l)P(\Delta_l=\delta_l)}\frac{P(O^l=j|\Delta_l=\delta_l)}{P(\tilde{O}^l=j|\Delta_l=\delta_l)} = \\
\frac{1}{8^{L-1}}\frac{P(O^l=j|\Delta_l=\delta_l,\Phi_l=\phi_l)P(\Delta_l=\delta_l,\Phi_l=\phi_l)}{P(O^l=j,\Delta_l=\delta_l)}\frac{P(O^l=j|\Delta_l=\delta_l)}{P(\tilde{O}^l=j|\Delta_l=\delta_l)} = \\
\frac{1}{8^{L-1}}P(\Phi_l=\phi_l|\Delta_l=\delta_l,O^l=j)\frac{P(O^l=j|\Delta_l=\delta_l)}{P(\tilde{O}^l=j|\Delta_l=\delta_l)} = \\
\frac{1}{8^L}\frac{P(O^l=j|\Delta_l=\delta_l)}{P(\tilde{O}^l=j|\Delta_l=\delta_l)}. \qquad (3.9)
\end{aligned}$$

Note that $P(O^l=j|\Delta_l=\delta_l)$ and $P(\tilde{O}^l=j|\Delta_l=\delta_l)$ are independent of ϕ_l, and hence

$$P\Big(\Phi=(\phi_1,...,\phi_L)\Big|O=j,\Delta=(\delta_1,...,\delta_L)\Big)$$

is also independent of ϕ_l. The same result holds for any $l=1,...,L$, hence the proof is completed. ■

Recall that a single X/Z Dango chain is used for the implementation of a fixed X/Z rotation. The following lemma shows this information does not help Bob to learn about Alice's secret.

Lemma 3. *Under the assumptions of Lemma 2, assume that Φ takes values with a non-zero probability only in a subset K, i.e. $\Phi\in K\subset\mathcal{A}^{\times L}$. Then*

$$P\Big(\Phi=\phi\Big|O=j,\Delta=(\delta_1,...,\delta_L),\Phi\in K\Big)=\frac{1}{|K|}$$

for any $\phi\in K$, $(\delta_1,...,\delta_L)\in\mathcal{A}^{\times L}$, $j\in\{1,...,m\}$, and POVM on $\bigotimes_{i=1}^L\rho_{\Theta_i}$.

Proof: Similar to the previous proofs, we have

$$P\Big(\Phi=\phi\Big|O=j,\Delta=(\delta_1,...,\delta_L),\Phi\in K\Big)=\frac{P\Big(\Phi=\phi,\Phi\in K\Big|O=j,\Delta=(\delta_1,...,\delta_L)\Big)}{P\Big(\Phi\in K\Big)}=\frac{P\Big(\Phi=\phi\Big|O=j,\Delta=(\delta_1,...,\delta_L)\Big)}{P\Big(\Phi\in K\Big)}=\frac{1/8^L}{|K|/8^L}.$$

■

Theorem 1. *The single-server protocol satisfies (S1).*

Proof:

Bob receives $2NM$ Dango states, and therefore there are $2NM$ secret angles $\{\theta_{a,b}^{Z/X}\}_{(a,b)=(1,1)}^{(N,M)}$. Bob also receives $2NM$ angles $\{\delta_{a,b}^{Z/X}\}_{(a,b)=(1,1)}^{(N,M)}$ from Alice. Let $\Phi,\Delta\in\mathcal{A}^{\times 2NM}$ be random variables, and $O\in\{1,...,m\}$ be the random variable which corresponds to the result of the POVM measurement which Bob performs on his system. Since Bob knows that Alice will attempt the same rotation many times in a single Z/X-AKLT subsystem until she succeeds, and that after the success of the rotation Alice will implements the trivial identity operation on the rest of qutrits in the Z/X-AKLT subsystem, Bob can assume that Φ takes values only in a subset K: $\Phi\in K\subset\mathcal{A}^{\times 2NM}$, where $|K|=8^{2NM/n}$.

From Lemma 1, 2, and 3, we have the following equality $\forall\phi\in K$, $\forall\delta_{a,b}^{Z/X}\in\mathcal{A}$, $(a=1,...,N)$, $(b=1,...,M)$, $j\in\{1,...,m\}$, and for any POVM

$$P\Big(\Phi=\phi\Big|O=j,\Delta=\{\delta_{a,b}^{Z/X}\}_{(a,b)=(1,1)}^{(N,M)},\Phi\in K\Big)=\frac{1}{|K|}$$

■

Theorem 2. *The single-server protocol satisfies (S2).*

Proof: It is easy to see

- When $R_Z^\alpha(\phi_{k,b}^Z,r_{(k-1)n+j,b}^Z)$ or $R_Z^\beta(\phi_{k,b}^Z,r_{(k-1)n+j,b}^Z)$ is implemented in the correlation space, the byproduct $XZ^{r_{(k-1)n+j,b}^Z}$ or $XZ^{r_{(k-1)n+j,b}^Z+1}$ occurs, respectively. If Bob has no information about the value of $r_{(k-1)n+j,b}^Z$, he cannot know whether the byproduct Z appears or not.
- When $R_X^\alpha(\phi_{k,b}^X,r_{(k-1)n+j,b}^X)$ or $R_X^\beta(\phi_{k,b}^X,r_{(k-1)n+j,b}^X)$ is implemented in the correlation space, the byproduct $X^{r_{(k-1)n+j,b}^X+1}Z$ or $X^{r_{(k-1)n+j,b}^X}Z$ occurs, respectively. If Bob has no information about the value of $r_{(k-1)n+j,b}^X$, he cannot know whether the byproduct X appears or not.

Note that we assume Alice's computation is implemented via a regular structure hence it contains both X and Z rotations. Therefore, both byproducts of Pauli X and Z operators will appear leading to the full one-time padding of the computation in the correlation space. In fact, we can show that Bob cannot have any information about the values of $\{r^{Z/X}_{(k-1)n+j,b}\}$ by showing similar proofs as those for $\{\phi^{Z/X}_{k,b}\}$. However, it is easy to consider that $\phi^{Z/X}_{k,b}$ takes values only 0 or π in the above proofs leading to the exchange of the role of $\{r^{Z/X}_{(k-1)n+j,b}\}$ and $\{\phi^{Z/X}_{k,b}\}$.

3.6.3 The two-server protocol

In this section, we will explain the two server protocol (Figure 3.3). There are two advantages behind the new protocol: Alice could be completely classical and more importantly for Bob(s) the resource state preparation and computation could be done more robustly using the energy-gap protection. To achieve these new features while keeping the security requirement intact, it is assumed that the two servers, Bob_1 and Bob_2, share many Bell pairs but have no classical or quantum channel between them. As we will discuss later it is an interesting open question (both from the practical and theoretical perspective) whether this assumption could be relaxed.

In the two server protocol, Bob_1 first creates AKLT resource states (without any random rotation), hence the preparation and storage of the state could be performed using ground state energy-gap protection as described in detail in [49]. Next, depending on Alice's desired gates, Bob_1 adiabatically turns off the interaction between some particles and the rest of particles in his resource state, and teleports the states of these particles to Bob_2 by consuming Bell pairs. Bob_1 sends Alice the result of the measurement in the teleportation through the classical channel. Note that due to the lack of any communication (classical or quantum) channels between Bob_1 and Bob_2, the teleportation procedure from Bobs' point of view can be seen as a usage of a totally mixed channel where only Alice knows how to correct the output of the channel.

Next, Alice calculates the angle in which particles should be measured by using her classical computer, and sends Bob_2 the angle that is the sum of the calculated angle plus the compensation for the byproduct and a random angle to hide the actual angle of the computation. Bob_2 performs the measurement in that angle and sends the result of the measurement to Alice. Next, Alice sends the previous random angle to Bob_1 and he does the single-qubit rotation which compensates the added random angle. Bob_1 and Bob_2 repeat this two-way classical communication with Alice until they finish the computation.

We define a (k,b)th AKLT subsystem ($k = 1, ..., N/n$, $b = 1, ..., M$) to be the collection of n qutrits of the bth AKLT chain with column index $(k-1)n+1, (k-1)n+2, ..., (k-1)n+n$ (Figure 3.4 (a)). A single-qubit rotation is implemented in a single AKLT subsystem. Let us assume that Alice wants to perform the single-qubit Z-rotation $\exp\left[\frac{iZ}{2}\phi^{Z}_{k,b}\right]$ with $\phi^{Z}_{k,b} \in \mathcal{A}$ using (k,b)th AKLT subsystem (11). The protocol for implementing an arbitrary X-rotation is analogous and omitted here. Finally, in order to perform blind $\wedge Z$ gates, similar to the single-sever protocol, Bob_1 periodically implements $\wedge Z$ gates. In order to keep the register state in the ground space, the interactions are adiabatically turned off before each $\wedge Z$ gate. Unwanted $\wedge Z$ gates are canceled in the same way as that in the single-server protocol.

Protocol 11 Two server Blind Z rotation

Initially the flag parameter (known to only Alice) is set $\tau = 1$. Alice sets her secret parameter $\omega^Z_{k,b} = 0$ and also chooses random angles $\xi^Z_{k,b} \in \mathcal{A}$. Alice sends Bob_1 parameter values N, M and $n < N$. Bob_1 creates M AKLT chains $|AKLT_b^{N,L,R}\rangle$, where $b = 1, \ldots, M$ (see equation 3.8) of N qutrits arranged in an array of N columns and M rows. For $j = 1 \cdots n/2$ Alice, Bob_1, and Bob_2 repeat (I)-(VI).

(I) Bob_1 adiabatically turns off the interaction, which acts on the jth qutrit of (k, b)th AKLT subsystem, and applies $P^\dagger$ to the isolated qutrit in order to convert the qutrit into the pair of two qubits (Figure 3.4 (b)). (The application of $P^\dagger$ can be done deterministically.)

(II) Bob_1 teleports the created two qubits to Bob_2 by consuming two Bell pairs (Figure 3.4 (c)). These two teleported qubits are affected by a two-qubit Pauli error $E \otimes E' \in \{I, X, Z, XZ\} \otimes \{I, X, Z, XZ\}$ (Figure 3.4 (d)). Bob_1 sends Alice the result of the measurement in the teleportation and hence only Alice and Bob_1 knows which error appears.

(III) Bob_2 applies the filtering operation $\{I - |\eta_1\rangle\langle\eta_1|, |\eta_1\rangle\langle\eta_1|\}$ to the received two qubits, and sends the result to Alice.

- If the Pauli error is $I \otimes I$, $X \otimes X$, $Z \otimes Z$, or $XZ \otimes XZ$, the probability of realizing $|\eta_1\rangle\langle\eta_1|$ is 0.
- If the Pauli error is $I \otimes Z$, $X \otimes XZ$, $Z \otimes I$, or $XZ \otimes X$, $|\eta_1\rangle\langle\eta_1|$ is realized with the probability $1/3$. If $|\eta_1\rangle\langle\eta_1|$ is realized, Z is implemented in the correlation space.
- If the Pauli error is $I \otimes X$, $X \otimes I$, $Z \otimes XZ$, or $XZ \otimes Z$, $|\eta_1\rangle\langle\eta_1|$ is realized with the probability $1/3$. If $|\eta_1\rangle\langle\eta_1|$ is realized, X is implemented in the correlation space.
- If the Pauli error is $I \otimes XZ$, $X \otimes Z$, $Z \otimes X$, or $XZ \otimes I$, $|\eta_1\rangle\langle\eta_1|$ is realized with the probability $1/3$. If $|\eta_1\rangle\langle\eta_1|$ is realized, XZ is implemented in the correlation space.

If $|\eta_1\rangle\langle\eta_1|$ is realized, skip steps (**IV**), (**V**) and (**VI**). If $I - |\eta_1\rangle\langle\eta_1|$ is realized, Bob_2 further applies the PEPS operation P on the two qubits. This PEPS operation is done deterministically, because the two qubits are already projected by $I - |\eta_1\rangle\langle\eta_1|$.

(IV) Alice sends the angle $\delta^Z_{(k-1)n+j,b}$ to Bob_2. This angle is determined according to the following rule:

- If the Pauli error is $I \otimes I$, $I \otimes Z$, $Z \otimes I$, or $Z \otimes Z$,

$$\delta^Z_{(k-1)n+j,b} = \tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)n+j,b} + r^Z_{(k-1)n+j,b}\pi \pmod{2\pi},$$

  where $\theta^Z_{(k-1)n+j,b} \in \mathcal{A}$ and $r^Z_{(k-1)n+j,b} \in \{0, 1\}$ are random numbers chosen by Alice, and signs of $\phi^Z_{k,b}$ and $\xi^Z_{k,b}$ should be changed if there is the byproduct X before this step.
- If the Pauli error is $X \otimes X$, $X \otimes XZ$, $XZ \otimes X$, or $XZ \otimes XZ$,

$$\delta^Z_{(k-1)n+j,b} = -\tau\phi^Z_{k,b} - \tau\xi^Z_{k,b} - \theta^Z_{(k-1)n+j,b} + r^Z_{(k-1)n+j,b}\pi \pmod{2\pi},$$

  where $\theta^Z_{(k-1)n+j,b} \in \mathcal{A}$ and $r^Z_{(k-1)n+j,b} \in \{0, 1\}$ are random numbers chosen by Alice, and signs of $\phi^Z_{k,b}$ and $\xi^Z_{k,b}$ should be changed if there is the byproduct X before this step.
- If the Pauli error is $I \otimes X$, $I \otimes XZ$, $X \otimes I$, $X \otimes Z$, $Z \otimes X$, $Z \otimes XZ$, $XZ \otimes I$, or $XZ \otimes Z$,

$$\delta^Z_{(k-1)n+j,b} = \theta^Z_{(k-1)n+j,b}$$

  where $\theta^Z_{(k-1)n+j,b} \in \mathcal{A}$ is a random number chosen by Alice.

(V) Bob_2 does the measurement $\mathcal{M}(\delta^Z_{(k-1)n+j,b})$ (similar to Protocol 2), and sends the result to Alice. By this measurement, following operations are implemented in the correlation space:

- If the Pauli error is $I \otimes I$ or $Z \otimes Z$, $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$, $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$, or $|\gamma\rangle$ occurs with the probability $1/3$ respectively. If $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$ is realized,

$$XZ^{r^Z_{(k-1)n+j,b}} \exp\Big[\frac{iZ}{2}\Big(\tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)n+j,b}\Big)\Big]$$

  is implemented. If $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$ is realized,

$$XZ^{r^Z_{(k-1)n+j,b}+1} \exp\Big[\frac{iZ}{2}\Big(\tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)+j,b}\Big)\Big]$$

  is implemented. If $|\gamma\rangle$ is realized, Z is implemented.

Protocol 11 — Continued

(V) —Continued

- If the Pauli error is $I \otimes Z$ or $Z \otimes I$, $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$ or $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$ occurs with the probability $1/2$ respectively. If $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$ is realized,

$$XZ^{r^Z_{(k-1)n+j,b}+1} \exp\Big[\frac{iZ}{2}\Big(\tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)n+j,b}\Big)\Big]$$

  is implemented. If $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$ is realized,

$$XZ^{r^Z_{(k-1)n+j,b}} \exp\Big[\frac{iZ}{2}\Big(\tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)n+j,b}\Big)\Big]$$

  is implemented.
- If the Pauli error is $X \otimes X$ or $XZ \otimes XZ$, $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$, $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$, or $|\gamma\rangle$ occurs with the probability $1/3$ respectively. If $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$ is realized,

$$XZ^{r^Z_{(k-1)n+j,b}} \exp\Big[\frac{iZ}{2}\Big(\tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)n+j,b}\Big)\Big]$$

  is implemented. If $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$ is realized,

$$XZ^{r^Z_{(k-1)n+j,b}+1} \exp\Big[\frac{iZ}{2}\Big(\tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)n+j,b}\Big)\Big]$$

  is implemented. If $|\gamma\rangle$ is realized, Z is implemented.
- If the Pauli error is $X \otimes XZ$ or $XZ \otimes X$, $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$ or $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$ occurs with the probability $1/2$ respectively. If $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$ is realized,

$$XZ^{r^Z_{(k-1)n+j,b}+1} \exp\Big[\frac{iZ}{2}\Big(\tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)n+j,b}\Big)\Big]$$

  is implemented. If $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$ is realized,

$$XZ^{r^Z_{(k-1)n+j,b}} \exp\Big[\frac{iZ}{2}\Big(\tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)n+j,b}\Big)\Big]$$

  is implemented.
- If the Pauli error is $I \otimes X$, $X \otimes I$, $Z \otimes XZ$, or $XZ \otimes Z$, $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$, $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$, or $|\gamma\rangle$ occurs with the probability $\frac{1}{2}\sin^2[\frac{1}{2}\delta^Z_{(k-1)n+j,b}]$, $\frac{1}{2}\cos^2[\frac{1}{2}\delta^Z_{(k-1)n+j,b}]$, or $1/2$, respectively. If $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$ or $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$ is realized, Z is implemented. If $|\gamma\rangle$ is realized, XZ is implemented.
- If the Pauli error is $I \otimes XZ$, $X \otimes Z$, $Z \otimes X$, or $XZ \otimes I$, $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$, $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$, or $|\gamma\rangle$ occurs with the probability $\frac{1}{2}\cos^2[\frac{1}{2}\delta^Z_{(k-1)n+j,b}]$, $\frac{1}{2}\sin^2[\frac{1}{2}\delta^Z_{(k-1)n+j,b}]$, or $1/2$, respectively. If $|\alpha(\delta^Z_{(k-1)n+j,b})\rangle$ or $|\beta(\delta^Z_{(k-1)n+j,b})\rangle$ is realized, Z is implemented. If $|\gamma\rangle$ is realized, X is implemented.

(VI) If the z-rotation $\exp[\frac{iZ}{2}(\tau\phi^Z_{k,b} + \tau\xi^Z_{k,b} + \theta^Z_{(k-1)n+j,b})]$ is implemented in the previous step, Alice sets $\tau = 0$, and $\omega^Z_{k,b} = \omega^Z_{k,b} + \theta^Z_{(k-1)n+j,b} \pmod{2\pi}$ (if there is no X byproduct before this rotation) or $\omega^Z_{k,b} = \omega^Z_{k,b} - \theta^Z_{(k-1)n+j,b} \pmod{2\pi}$ (if there is the X byproduct before this rotation).

(VII) So far, the z-rotation

$$G^Z_{k,b} \exp\Big[\frac{iZ}{2}(\omega^Z_{k,b} + \xi^Z_{k,b})\Big] \exp\Big[\frac{iZ}{2}\phi^Z_{k,b}\Big]$$

up to some Pauli byproduct $G^Z_{k,b}$ is implemented. The probability that they fail to perform this z-rotation is $(2/3)^{n/2}$ can be made arbitrarily small by choosing a sufficiently large n. Alice asks Bob_1 to correct the accumulated error. In order to do so she sends Bob_1 the angle $\tilde{\omega}^Z_{k,b} = \omega^Z_{k,b} + \xi^Z_{k,b} \pmod{2\pi}$ if $G^Z_{k,b}$ contains no X byproduct, and $\tilde{\omega}^Z_{k,b} = -\omega^Z_{k,b} - \xi^Z_{k,b} \pmod{2\pi}$ if $G^Z_{k,b}$ contains the X byproduct. Bob_1 implements the rotation $\exp\Big[-\frac{iZ}{2}\tilde{\omega}^Z_{k,b}\Big]$ by using the rest of the qutrits in (k,b)th AKLT subsystem. The probability that Bob_1 fails to perform this z-rotation is $(1/3)^{n/2}$, which can be made arbitrarily small by choosing a sufficiently large n.

3.6.4 Proof of blindness of the two server protocol

In this section, we will show the blindness of the two server protocol.

Definition 2. *A two sever protocol is blind if*

(D1) The conditional probability distribution of Alice's nontrivial computational angles, given all the classical information Bob_1 can obtain during the protocol, and given the measurement results of any POVMs which Bob_1 may perform on his system at any stage of the protocol, is uniform,

(D2) The conditional probability distribution of Alice's nontrivial computational angles, given all the classical information Bob_2 can obtain during the protocol, and given the measurement results of any POVMs which Bob_2 may perform on his system at any stage of the protocol, is uniform,

(D3) The register state of Bob_1 in the correlation space is quantum one-time padded,

(D4) The register state of Bob_2 in the correlation space is quantum one-time padded.

The proof is based on following lemmas.

Lemma 4. *Bob_1 cannot send any information to Bob_2.*

Proof: By the assumption, there is no channel between Bob_1 and Bob_2. Furthermore, Bob_1 cannot send any information to Bob_2 via Alice either. This is due to the following facts. First, what Bob_1 sends to Alice are the "results of measurements in teleportations". Second, what Alice sends to Bob_2 are $\{\delta^{Z/X}_{(k-1)n+j,b}\}$. Third, recall that Alice chooses the definition of each $\delta^{Z/X}_{(k-1)n+j,b}$ among

$$\begin{aligned}\delta^{Z/X}_{(k-1)n+j,b} &= \tau\phi^{Z/X}_{k,b} + \tau\xi^{Z/X}_{k,b} + \theta^{Z/X}_{(k-1)n+j,b} + r^{Z/X}_{(k-1)n+j,b}\pi \pmod{2\pi},\\ \delta^{Z/X}_{(k-1)n+j,b} &= -\tau\phi^{Z/X}_{k,b} - \tau\xi^{Z/X}_{k,b} - \theta^{Z/X}_{(k-1)n+j,b} + r^{Z/X}_{(k-1)n+j,b}\pi \pmod{2\pi},\end{aligned}$$

or

$$\delta^{Z/X}_{(k-1)n+j,b} = \theta^{Z/X}_{(k-1)n+j,b}$$

according to what Bob_1 sends to Alice. However, the value of each $\delta^{Z/X}_{(k-1)n+j,b}$ is independent of what Bob_1 sends to Alice, since $\theta^{Z/X}_{(k-1)n+j,b}$ is completely random and therefore $\delta^{Z/X}_{(k-1)n+j,b}$ takes any value in $\mathcal{A}$ with equal probability, whichever definition Alice chooses. Therefore,

$$\begin{aligned}P\Big(T=t\Big|\Delta=\{\delta^{Z/X}_{(k-1)n+j,b}\}\Big) &= \frac{P\Big(\Delta=\{\delta^{Z/X}_{(k-1)n+j,b}\}\Big|T=t\Big)P\Big(T=t\Big)}{P\Big(\Delta=\{\delta^{Z/X}_{(k-1)n+j,b}\}\Big)}\\ &= P\Big(T=t\Big),\end{aligned}$$

where

T is the random variable which represents teleportation results.

Finally, just by using the Bell pairs shared between Bob_1 and Bob_2 no information can be sent from Bob_1 to Bob_2 without sending a classical message, as this would imply no-signaling. ■

Lemma 5. *Bob_2 cannot send any information to Bob_1.*

Proof: Similar to the previous lemma, Bob_2 cannot send any information to Bob_1 via Alice. This is due to the following facts. First, what Bob_2 sends to Alice are the "results of filterings and measurements". Second, what Alice sends to Bob_1 are $\{\tilde{\omega}_{k,b}^{Z/X}\}$. Third, although $\omega_{k,b}^{Z/X}$ depends on what Bob_2 sends Alice, $\tilde{\omega}_{k,b}^{Z/X}$ is independent of what Bob_2 sends Alice, since $\xi_{k,b}^{Z/X}$ is completely random. Therefore,

$$\begin{aligned} P\Big(F=f\Big|E=\{\tilde{\omega}_{k,b}^{Z/X}\}\Big) &= \frac{P\Big(E=\{\tilde{\omega}_{k,b}^{Z/X}\}\Big|F=f\Big)P\Big(F=f\Big)}{P\Big(E=\{\tilde{\omega}_{k,b}^{Z/X}\}\Big)} \\ &= P\Big(F=f\Big), \end{aligned}$$

where F is the random variable which represents the results of filterings and measurements. ■

Theorem 3. *The two server protocol satisfies (D2).*

Proof: From Lemma 4, Bob_1 cannot send any information to Bob_2. Therefore, all quantum states which Bob_2 receives are completely mixed states, and all classical information which Bob_2 obtains are only the angles $\{\delta_{(k-1)n+j,b}^{Z/X}\}$. Since $\{\theta_{(k-1)n+j,b}^{Z/X}\}$ and $\{\xi_{k,b}^{Z/X}\}$ are completely random and independent from $\{\phi_{k,b}^{Z/X}\}$, Bob_2 cannot have any information about $\{\phi_{k,b}^{Z/X}\}$ from $\{\delta_{(k-1)n+j,b}^{Z/X}\}$. ■

Theorem 4. *The two server protocol satisfies (D4).*

Proof: It is easy to see that

- When the z-rotation $\exp\left[\frac{iZ}{2}\Big(\phi_{k,b}^{Z}+\xi_{k,b}^{Z}+\theta_{(k-1)n+j,b}^{Z}\Big)\right]$ is implemented, the byproduct $XZ^{r_{(k-1)n+j,b}^{Z}}$ or $XZ^{r_{(k-1)n+j,b}^{Z}+1}$ occurs.
- When the x-rotation $\exp\left[\frac{-iX}{2}\Big(\phi_{k,b}^{X}+\xi_{k,b}^{X}+\theta_{(k-1)n+j,b}^{X}\Big)\right]$ is implemented, the byproduct $X^{r_{(k-1)n+j,b}^{X}}Z$ or $X^{r_{(k-1)n+j,b}^{X}+1}Z$ occurs.

Bob_2 cannot gain any information about $\{r_{(k-1)n+j,b}^{Z/X}\}$ from $\{\delta_{(k-1)n+j,b}^{Z/X}\}$, since $\{\theta_{(k-1)n+j,b}^{Z/X}\}$ and $\{\xi_{k,b}^{Z/X}\}$ are completely random and independent from $\{r_{(k-1)n+j,b}^{Z/X}\}$.

Theorem 5. *The two-server protocol satisfies (D3).*

Proof: First, since $\{\theta_{(k-1)n+j,b}^{Z/X}\}$ are completely random, hence $\{\omega_{k,b}^{Z/X}\}$ is independent from $\{r_{(k-1)n+j,b}^{Z/X}\}$. Second, although $\tilde{\omega}_{k,b}^{Z/X}$ is related to the parity of X or Z in $G_{k,b}^{Z/X}$, Bob_1 cannot know these parities from $\{\tilde{\omega}_{k,b}^{Z/X}\}$ since Bob_1 does not know $\{\xi_{k,b}^{Z/X}\}$. Therefore (D3) is satisfied. ■

Theorem 6. *The two-server protocol satisfies (D1).*

Proof: From Lemma 5, Bob_2 cannot send any information to Bob_1. Therefore, the classical information which Bob_1 can obtain are only $\{\tilde{\omega}_{k,b}^{Z/X}\}$. However, $\{\tilde{\omega}_{k,b}^{Z/X}\}$ are independent from $\{\phi_{k,b}^{Z/X}\}$. Furthermore, from Theorem 5, Bob_1's states are one-time padded to him, therefore, no POVM on Bob_1's states gives information to Bob_1. ■

3.6.5 Span of encrypted AKLT states

Here we calculate the span of all pre-rotated AKLT states.

$$\left\{ |RAKLT_b^{2N,L,R}(\{\theta_{a,b}^{Z/X}\})\rangle \;\middle|\; \theta_{a,b}^{Z/X} \in \mathcal{A},\ a = 1, ..., N \right\}$$

As is shown below, the dimension of the span is 2^{2N}. Therefore, if Bob does not know $\{\theta_{a,b}^{Z/X}\}$, he must prepare an unnatural Hamiltonian with exponentially-degenerated ground states.

Let us show that the dimension of the span is 2^{2N}. Let

$$|\psi_{2N}\rangle = |\delta_0\rangle \otimes \Big(\bigotimes_{i=1}^{2N} |\delta_i\rangle \Big) \otimes |\delta_{2N+1}\rangle,$$

where $|\delta_0\rangle$ and $|\delta_{2N+1}\rangle$ are qubit states, and $|\delta_i\rangle$ ($i = 1, ..., 2N$) are qutrit states.

Let

$$\mathbf{U}_\theta = I_2 \otimes \mathcal{U}_b(\{\theta_{a,b}^{Z/X}\}) \otimes I_2$$

be a global unitary operator, where I_2 is the identity operator on a single qubit. The parametrized unitary map $\mathbf{U}$ is the pre-rotation map, which when applied on an AKLT state generates the encrypted resource state Alice will have Bob use in a run of a UBQC protocol, parametrized by Alice's secret angles $\theta_{a,b}^{Z/X}$. Let E be a global unitary operator which works as

$$E|\psi_{2N}\rangle = |AKLT^{2N,L,R}\rangle.$$

From Lemma below and the fact that $E^\dagger$ is unitary,

$$\begin{aligned}
\dim\Big(\text{span}\Big\{\mathbf{U}_\theta E|\psi_{2N}\rangle\Big\}_\theta\Big) &= \dim\Big(\text{span}\Big\{(\mathbf{U}_\theta E)^\dagger|\psi_{2N}\rangle\Big\}_\theta\Big) \\
&= \dim\Big(\text{span}\Big\{E^\dagger\mathbf{U}_\theta^\dagger|\psi_{2N}\rangle\Big\}_\theta\Big) \\
&= \dim\Big(\text{span}\Big\{\mathbf{U}_\theta^\dagger|\psi_{2N}\rangle\Big\}_\theta\Big) \\
&= \dim\Big(\text{span}\Big\{\mathbf{U}_\theta|\psi_{2N}\rangle\Big\}_\theta\Big) \\
&= 2^{2N}.
\end{aligned}$$

Lemma: Let $\{V_1, \dots V_r\}$ be a set of r operators, and let $|\delta\rangle$ be a state in their domain. Then

$$\dim\Big(\text{span}\Big\{V_i|\delta\rangle\Big\}_i\Big) = \dim\Big(\text{span}\Big\{V_i^\dagger|\delta\rangle\Big\}_i\Big).$$

Proof: Recall that $\dim\Big(\text{span}\Big\{V_i|\delta\rangle\Big\}_i\Big)$ is equal to the rank of the Gram matrix of the set of vectors $\{V_i|\delta\rangle\}_i$. Also note that if G_A is the Gram matrix of the set of vectors $\{V_i|\delta\rangle\}_i$ and G_B is the Gram matrix of the set of vectors $\{V_i^\dagger|\delta\rangle\}_i$, then $G_A = G_B^*$. Finally, let us remind that $\text{rank}(A) = \text{rank}(A^*)$ for all matrices A.

Chapter 4

Discussion: future of UBQC

Here we discuss some of the directions in which UBQC may develop

Throughout the preceding chapters we have addressed the basic ideas behind UBQC. We have focused on the privacy guaranteed to the client, and discussed two aspects of this protocol as far as real-life realizability of it is concerned. However, many other aspects of this protocol and its variations, the development of which occurred in parallel to the work presented, we have left unmentioned. We remedy this lapse in this chapter by looking into the cutting edge and the future of UBQC.

4.1 Verifiable UBQC

The two main properties we focused on, which we require from any useful delegated computation scheme are **correctness** and **blindness**. The correctness property entailed that, provided both the server and the client acted *honestly*, meaning according to the prescription of the protocol, the output of the computation was indeed correct. We did not have to pay overly detailed attention to correctness. Because of the relatively straightforward nature of how the computation is encrypted in UBQC, the correctness followed almost directly from the correctness of computation in the underlying models we considered: the one-way model and generalized MBQC on AKLT states. However, the second property, blindness, required close attention. Here, we focus on a third property of great interest.

In the setting of delegated computation, the property of **verifiability** would be a great asset. Roughly speaking, a delegated computation protocol is verifiable, if there exists a mechanism which guarantees to the client that the output of the computation the client derives from the server is indeed correct. Recall, in UBQC the client is assumed to be comparatively less powerful than the server. In particular, the client is assumed to have the technological capabilities which can easily be achieved in practice today - a classical computer, a source of randomness of a suitable quality [1] and a generator of single qubit states $\{|+_\theta\rangle\}_\theta$ (see 1), realized for instance by a single-

[1] The issue of "quality" of randomness, and the value of it as a resource is an important and subtle field of research. For simplicity, we shall take a naive approach and assume the client has access to a fair coin without

photon source and a controlled polarizer.

For certain classes of computational problems, the client can check the final outcome of the delegated computation received from the server directly. For instance, assume the client delegated an integer factoring problem to the server. This problem is efficiently solvable on a quantum computer using the famous Shor's algorithm, but not known to be efficiently solvable on a classical computer. Nonetheless, the client can easily check whether the solution is correct, by division, which is efficient on classical machines. Similar holds for other problems known to be efficiently solvable on quantum machines. In complexity theory, this property of a problem – that the solution is easily checked – characterises a class known as NP. However, one of the most interesting features of quantum computers is that it seems that some of the problems they can solve efficiently do not belong to the class NP – the problem of estimating the value the Jones polynomial attains at particular chosen points is not believed to be in NP and yet, is solvable on a quantum computer. Formally, it is believed that $\mathsf{BQP} \not\subseteq \mathsf{NP}$, where the class of problems BQP are those solvable efficiently by a quantum computer [2]. The Jones polynomial problem is also *complete* for the BQP class, meaning every instance of a problem in BQP can efficiently be reduced to an instance of the Jones polynomial problem of the same size. Thus, if the client has the need to BQP−hard problems using UBQC, it is likely she will have no means to verify the solution. More generally, UBQC needs not be used for decision problem solving alone. In the most general variant, the output of the computation is a quantum state, designed by the client using the servers system – a remote quantum laboratory! The experimentally-challenged client then receives a quantum state from the server, and verifying whether the state is what client wished to obtain may be difficult.

As we have mentioned in the introduction to Part 1 of this thesis, the notion of verifiable quantum computation was initially addressed by Arrighi and Salvail. Their solution was based on the idea of surrounding the actual relevant computation in a large number of "trap" computations. In order to detect a malicious behaviour from the server, the client is required to know the correct output of the computation. This restricted the class of functions which can be solved by using this protocol to the class of *random verifiable* functions as given in definition 1. The exact relationship of random-verifiable problems to BQP is unknown, however it is not believed they contain the entirety of BQP.

The direct application of the technique of "trap computations" to the UBQC scenario, where all the computations can be concealed from the server yields only moderate benefits. Since the computed function is not public in UBQC, there is now nothing preventing the client from using random-verifiable problems for traps, and any BQP problem for her actual computation. Nonetheless, it still may be the case that the server can only solve random-verifiable problems correctly. He will then not get caught, and the computational outcome for the actual non-random-verifiable problem may still be incorrect. Instead of directly applying this idea to UBQC, we gain more by adopting the basic concepts behind it: the only thing the client can do, to verify what the

being very precise about what this means. What we demand from the coin is that the variables which are used in UBQC, based on the result of a coin toss (or tosses), are indeed distributed uniformly at random, and uncorrelated to whatever system the adversary (or the environment) may have.

[2] It is also believed that $\mathsf{NP} \not\subseteq \mathsf{BQP}$.

server is actually doing, is to present to him a problem the client already knows the resolution to. In the setting of UBQC, the tasks/problems the client presents, which ensure the correctness of the computation are: 1) entanglement operations ($\wedge Z$) on pre-rotated qubit, and 2) measurements of individual qubits. Thus, the minimal process we can use to verify the servers behaviour is a few qubit processing task, checking 1) and 2) simultaneously. This is achieved by using the idea of trap qubits, initially introduced in [1].

To explain the basic idea, assume the client can generate the $|+_\theta\rangle$ states used in UBQC, but additionally, the states $|0\rangle$ and $|1\rangle$. Imagine a sufficiently large pre-rotated graph state (in the sense used in UBQC), generated from the qubits the server received from the client. Consider a particular qubit we denote T (for "trap"), somewhere in the middle of this resource state, pre-rotated to the state $|+_\theta\rangle$. The neighbours of T qubit are pre-set to one of the $\{|0\rangle, |1\rangle\}$ states chosen at random. Then, the control-Z interaction will not entangle the qubit T with the rest of the resource state – it is isolated. The state of the qubit, if the control-Z was indeed applied will be $|+_{\theta+p\pi}\rangle$ where p is the parity of the number of neighbours of the T qubit which were pre-set to the $|1\rangle$ state. The state of this qubit is known to the client alone, and if the client, during the run time of the actual computation, asks the server to measure T with respect to the angle θ, the outcome is deterministic and known to the client. In order for the server to report the correct outcome of the measurement on the trap qubit with unit probability, the server would have to apply the required $\wedge Z$ interactions, and measure the qubit with the correct angle. In principle, the client can position the trap anywhere in the resource state. This position is not known to the server, so in order to avoid reporting an incorrect outcome on the trap, with unit probability, the server has to perform all $\wedge Z$ operations and all the measurements correctly, ensuring correct computation. This idea general idea we will refer to as *trapification* illustrated in Figure 4.1.

To amplify the detection probability of a malicious behaviour of the server, two techniques can be employed. Firstly, one ensures that the number of trap qubits is substantially large, for instance proportional to the size of computation, while still allowing arbitrary computation. Secondly, the entire computation should be embedded in a fault-tolerant code which can handle a sufficient error weight. Then, intuitively, the amount of deviation the server may induce to the computation, provided he behaves correctly enough to pass the testing on the traps, will be corrected by the fault tolerant code. This guarantees the final correctness of the output, provided the server reported the correct outcome on all traps.

An explicit construction of a protocol which uses the idea of trapification invented by two of the original authors of [1] which satisfies all the issues addressed above is presented in [42].

The framework of approximate blindness, and the idea of remote blind qubit state preparation (RBSP) we have developed and presented in Chapter 2 could be compatible with the notion of verifiability. We emphasize that the arguments which follow are not claimed to be a proofs, but conjectures. Consider a UBQC protocol which is δ-verifiable, meaning that the outcome is correct, except with probability δ. Assume that the client's preparation phase was implemented using a perhaps more generalized variant of RBSP which guaranteed that the real implementation of the UBQC protocol is ϵ-blind. Then, it may be possible to show that the resulting protocol

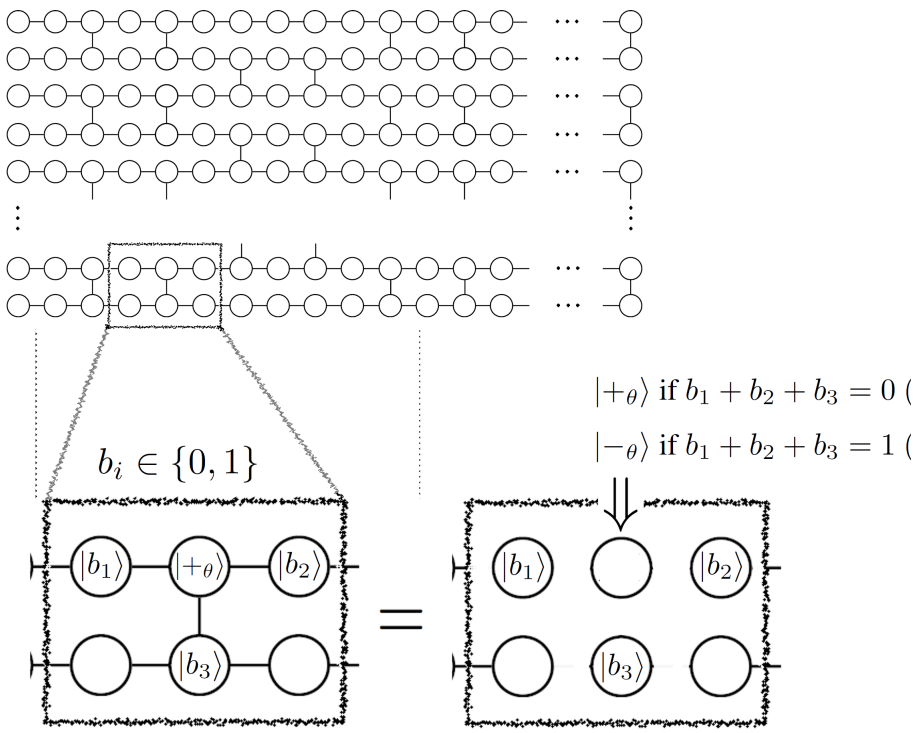

Figure 4.1. Trapification. If a graph-like state is built from individual qubits by using the $\wedge Z$ interaction, and a particular qubit T in the $|+_\theta\rangle$ state has neighbours in the Pauli-Z eigenstates, then the entangling operation will not entangle the qubit T to the rest of the system. The post-entangling (pure) state of the qubit depends on the initial pre-rotation ***and*** *the parity of the neighbours in the $|1\rangle$ state. Thus it can be used to check both whether the measurement and $\wedge Z$ interactions have been performed. The state of these qubits is known to Alice alone, and they can be used to check the behaviour of an untrusted server. The position of the traps can be hidden from the server because the server cannot distinguish between the trap-inducing qubit states ($1/2(|0\rangle\langle 0| + |1\rangle\langle 1|) = \mathbb{1}$) and randomly pre-rotated qubits used for the computation ($1/8 \sum_\theta |+_\theta\rangle\langle +_\theta| = \mathbb{1}$).*

is at least $\epsilon + \delta$-verifiable (and still ϵ-blind). This should hold as the measure of approximate blindness we use is based on the trace distance. This guarantees that whatever identical physical processes one performs on the ideal and real systems, the resulting quantum states, or realized probability distributions if measurements are involved, will not differ by more than ϵ in terms of the trace distance. Since the process of verification is a physical process done on the quantum systems the client and the server share, it should hold all resulting probabilities (including the δ in the ideal and δ' in the real protocol) should not differ by more than ϵ. The actual formal proof based on this observation we leave for the very immediate future. Verification itself has no direct influence on blindness in the case of UBQC (unlike in the case of the protocol in [18]). As we have seen, blindness holds even when verification is not present. However, it may be the case that verification (of the types suggested in [42], and [19]) cannot work without blindness. This is reminiscent of the fact that, in the quantum variants, message authentication necessarily implies message privacy [17].

Next, we consider the connections between a verifiable UBQC protocol and complexity theory.

UBQC and complexity theory UBQC is a game in which a powerful server solves problems for a more restricted client. Thus, it is natural to investigate the relationship between UBQC and the so-called *interactive proof systems*. Interactive proof systems are a category of abstract machines defined in a two-party setting. One imagines a restricted party called *verifier* and a computationally unbounded *prover*. The broad idea is as follows: to this pair a computational decision problem is presented and through interaction with the omnipotent prover, the verifier is required to correctly solve the problem. The problem is decidable in this class if two criteria are met:

- completeness – if the true decision problem solution is "YES", the prover will manage to convince the verifier of this fact with a probability above $p_c = 2/3$ [3], and
- soundness – if the true solution is "NO", a (corrupted) prover will succeed to convince the verifier of the opposite (the verifier will declare "YES") with probability no higher than $p_s = 1/3$..

Many variants of these classes exist, in particular, the power of the verifier may vary, but usually it is set to BPP. The allowed number of rounds of communication between the prover and the client may also vary from a constant, to polynomial. In the case of a BPP verifier, and a up to a polynomial number of communication rounds, the class of problems decidable by this abstract machine is called IP. Note that the class NP – the class of problems the solution for which can efficiently be verified – is the simplest non-trivial example of an IP class. Here, the prover simply sends the solution and a short proof to the verifier [4].

The UBQC protocol with verification very much resembles an IP setting: we have a powerful server (taking the part of the prover) and a restricted verifier. In UBQC, computationally the verifier is assumed to be a BPP just like in IP systems, however she is empowered to generate single qubit states. While such a capability is not within the domain of discourse of quantum complexity theory, we shall acknowledge this qubit-preparation power by stating that the verifier is BPP^* powerful. Strictly speaking, to run a verifiable UBQC protocol, the client requires less computational power than BPP. In particular, she needs only to perform $mod\ 8$ addition only in run-time [5]. The other important difference is that the prover is no longer omnipotent – he is restricted to BQP. The verifiable variant of the UBQC protocol empowers the client to

[3] For more information on the reasoning behind such bounded-error decision processes, refer to the definition of the BPP class at the end of this thesis.

[4] Strictly speaking, classes like NP and BQP are classes of decision problems, whereas, for instance, factoring is a functional problem. By means of *polynomial many-to-one reductions* it is possible to reduce many functional problems to decision problems. The class of decision problems NP can be verified efficiently in the following sense. For each instance size of a decision problem D in NP there exists an efficient algorithm V taking as input: 1) the decision $d = \{YES, NO\}$ of the problem instance, and 2) a *witness* W_d (a short proof). Then the efficient algorithm V, upon the input (d, W_d) returns YES if and only if d is the correct solution to the original problem D.

[5] The client may need to compute the structure of the fault tolerant code which is more difficult, but this computation is done off-line.

correctly and soundly decide BQP problems, which we can state in complexity theoretical terms as follows:

the class of problems BQP *admits an interactive proof system with a* BPP^* *verifier and a* BQP *prover.*

Such classical verifier–quantum prover IP systems were considered by Aharonov, Eban, Ben-Or in [19], and named QPIP (quantum prover interactive proof systems). As mentioned in Chapter 1, they also achieve verifiable blind quantum computing, but the client is assumed to have a constant-size universal quantum computer. A natural question which immediately comes to mind is the following: does BQP allow an interactive proof system with a strictly BPP client? As we will see, any answer to this question may have important consequences in physics. This question, paraphrased in the setting in UBQC asks: can we do UBQC with a completely classical Alice. The general answer to this question is unknown and seemingly difficult. It is the opinion of the author of this thesis that *if* such a protocol is possible it will be based on techniques not directly related to UBQC – as as we have seen in 3.3.1, the central idea to UBQC is the notion of "encrypting the (quantum) resource state", and privacy here is only maintained due to the specific nature of quantum states. A crucial role in blindness is played by the discrepancy between how much information is encoded in a quantum state which can be used, and how much classical information can be extracted from it. Recall, the angles θ_i encode 3 bits, and they are placed in the states $|+_\theta\rangle$ by the client, and all three bits influence the quantum computation which is performed. However, by the Holevo bound, at most 1 bit, more concretely, only the bit pertaining to the most significant digit in the binary encoding of θ can be extracted by the server (for details see 4.3.2). Classical systems do not have such curious properties, thus the encoding used in a fully classical client UBQC scheme might have to be something different. However, a fully classical client *is* achievable – in the two server setting. In this case we trade the client's "quantumness" for an additional rule that the two servers should not communicate. That places the two server solution outside the scope we are currently considering.

A UBQC protocol with a completely classical client could be linked to other known results in the classical computer science. In Chapter 1, we have mentioned that Feigenbaum et al. [23] have shown that no $\mathsf{NP} - hard$ function can be evaluated in an encrypted way, while leaking nothing but the instance size to the server by a ZPP client, using polynomially many rounds of communication, unless the polynomial hierarchy PH collapses to the third level. [6] If UBQC could be modified to work with a completely classical client, even without verification, this cold be used to prove $\mathsf{NP} \not\subseteq \mathsf{BQP}$, unless PH collapses to the third level. This would be a huge result on its own [78], but also because the proof of the relationship of this classes actively uses fundamentally cryptographic notions.

The fact that UBQC may have interesting consequences in complexity theory has already been established. Using the protocol with two servers, a new interesting result in the setting of multi-prover interactive proof systems was derived – it was used that prove $\mathsf{QMIP} = \mathsf{MIP}^*$ [79]. In multi-prover interactive proof systems (MIP), the verifier communicates with many computa-

[6] In their work, Feigenbaum et al. they consider the class ZPP as the class of effectively solvable problems. This class is defined in Section 10.3.2, and is contained in BPP.

tionally unbounded servers in order to decide the given problem. The servers are not allowed to communicate amongst themselves, otherwise there would be no difference between IP and MIP. To give a little bit of intuition why MIP should be more powerful than IP consider the following: the only reason the client cannot decide any difficult problem in an interactive proof system (more difficult than PSPACE) is because soundness, rather than completeness, cannot be ensured. Note that the prover always *knows* the result of the decision problem – he is computationally unbounded. But, this power also allows the prover to fool the verifier, and make the verifier accept an incorrect solution, if the problem is in to difficult a class. In the MIP setting, similarly to how in TV shows (maybe even in reality?) the inspectors take two suspects and interrogate them in separate rooms, to play one against the other, the verifier can do the same in an MIP setting. Indeed, the class MIP equals NEXPTIME which is far more powerful than IP $=$ PSPACE.

In the class QMIP, the verifier is a BQP machine, and in MIP*, the provers are assumed to share a sufficient amount of entanglement, *i.e.* a number of Bell pairs. Before the result in [79], it was known that MIP* $\subseteq$ QMIP. In [79] the authors use the idea of two-server UBQC, to empower a classical BPP verifier to perform BQP computation with two of the multi-provers, thereby effectively raising the verifier's powers from BPP to BQP. Now the claimed result seems obvious, since the only advantage QMIP may have over MIP* was in the lower powers of the verifier in the two classes.

While the claim of the innovative approach above is almost assuredly true, to be completely rigorous (even knit-picky), we should acknowledge that in [79] UBQC is used as a subroutine, called many times. UBQC was only proven to be secure as a stand-alone application, and not in any arbitrary setting. There are two possible resolutions to this problem: either one should present the proof of the claim MIP* $=$ QMIP without evoking the stand-alone blindness of UBQC, or, more elegantly, a compositional definition for UBQC must be derived, and a universally compositional security proven. We will return to the issue of universal composability presently. Next, we address the claim that proving BQP allows an interactive proof system with a (strictly) BPP client may have consequences on our understanding of (quantum) physics.

UBQC and physical reality Quantum systems seem difficult to simulate on a classical computer. The upside is, whatever it is that makes the simulation of quantum systems hard, can be used to build computers more powerful than the ones we have today. Provided simulation of quantum systems is difficult (equivalently, if BPP $\subsetneq$ BQP), then any attempt to verify, or more correctly, attempt to falsify the validity of quantum mechanics on large systems is just as difficult. If we cannot compute what the state of the system as predicted by the quantum theory should be, we certainly cannot verify whether what we get in the lab is consistent with it. This problem was pointed out by Vazirani in 2007. [80]. The seemingly purely complexity-theoretical question, whether a classical verifier can decide quantum problems with the aid of a quantum prover can be viewed as a formalization of the problem stated by Vazirani. UBQC, but also the results of [19], give a partial resolution: while we still know whether we can verify the behaviour of a large-scale quantum system (the server) with a completely classical device, we can do it

by manipulating smaller systems. In the case of UBQC this is achieved by manipulating only individual qubits, and in the case of in the case of [19] by using a constant size quantum computer. Additional prospects for interesting interplay between fundamental questions in quantum mechanics arise when one views UBQC as "renting a delegated quantum laboratory". While this may seem appealing for the financially-limited experimentalist, perhaps even more interesting features await discovery when the fact that the experiment performed in the lab is concealed from the lab itself! However, it is true that the security claims in UBQC rest on the basis of quantum mechanics, so verifying quantum mechanics, by initially assuming quantum mechanics may be without sense. We leave it as an open question if in this "blinded lab" setting interesting results may be achieved. For instance, could it be possible to achieve improvements with respect to the free will or signalling loopholes in Bell inequality violation experiments if we know the lab itself does not know it is doing a Bell test?

Mixing quantum physics and complexity theory has to be done carefully. Assume the server is owned by a genius who invented an algorithm which can simulate a quantum computer efficiently. Thus, he can solve BQP problems. Can this server pass the verification test in UBQC? Clearly no! The server is a classical machine, and to pass the verification test he would have to work with quantum states. One of the reasons this scenario occurs is because of the "$*$" in the BPP^* with which we denote the powers of the client in UBQC. "Qubits" are not in the domain of discourse when one discusses complexity classes for decision problems. While a BQP oracle can decide all the decision problems a quantum computer can (it is its definition, in fact), it clearly cannot manipulate quantum states, generate a qubit and so on. To illustrate with a more general example how the semantic confusion we are facing can produce funny sounding sentences consider the following imaginary conversation. Turing : "Ha! I have invented a Universal (Turing) Machine!" Hungry guy: "Universal? Great! Have it make me a sandwich!". This then brings the next obvious question: if we are not checking the computational power of the server, what are we checking? What does a passed verification of this form entail (beyond proving the server has BQP computational powers)? Understanding the answers to these questions will clarify the role UBQC may hold in understanding quantum mechanics itself better.

If UBQC with verification could be performed with a completely classical client, then the issue of what we actually are verifying could be relevant to the problems which arise when quantum mechanics gets involved in settings where many protocols communicate – a quantum cloud computing setting [81].

There are at least three types of "verification" we may be interested in which we now very briefly present.

Verification 1 (UBQC-type verification) This is the type of verification achieved in the single server UBQC protocol with a qubit-generation-capable client.

Verification 2 (complexity-theoretical) This type of verification is compatible with the complexity-theoretical notion of verification in interactive proof systems. In this setting, a decision problem

belonging to a particular complexity class (BQP) is set before the client (verifier), and the problem is verifiable if the client can correctly decide the problem by querying the server a certain number of times. A complexity class is verifiable in this fashion if all problems in it are verifiable. The server is assumed to be malicious.

Verification 3 (quantum mechanical) Here, the client verifies whether the server is a quantum machine. In particular, we may be interested in distinguishing between a (perhaps unrealistically fast) classical computer and a full fledged quantum computer.

We now briefly explore the relationships between the first verification type and the latter two and observe whether the tests could be performed by a completely classical client.

Verification 1 vs. Verification 2 In the UBQC verification, the client effectively observes whether the server performs the "quantum" computational commands, administered by the client, to such an extent that a fault-tolerant code can still distil a correct computational outcome from the process. This means that Verification 1 implies Verification 2. However there is an important distinction from the two. In Verification 2, it is clear that a client can verify the solution to the problem she is asked to decide, if the problem is such that she can solve it on her own. More generally, a P (or BPP) client can verify any NP problem by running an efficient verification procedure on the solution proposed by the server. However, then *the verification procedure depends on the problem being solved*. In the UBQC-type verification, the procedure for verification is *independent* from the problem being solved – in particular, the client need not even know what problem is being solved, she just needs to observe the behaviour of the server on the trap qubits. This distinction could also be stated in complexity-theoretical terms. To perform verification (ignoring the off-line computation required for the construction of the fault tolerant encoding) the client only needs to run $\mathrm{mod}\ 8$ arithmetic and a bit of binary comparison. This is a weaker requirement than to, say preform long division, required to check whether a number reported by the server really divides the input integer in a run of a factoring problem.

Verification 1 vs. Verification 3 Clearly, in Verification 1 the client checks whether *something* quantum is happening at the server's side – she sends him quantum states, and a purely classical machine would have no interface to do anything with them. But, as we have seen, the client actually checks a lot more – whether a particular POVM has been implemented by the server to a sufficient precision. Thus, a client capable of Verification 1 can perform Verification 3. Verification 3 could be performed using a much simpler protocol. The client prepares and sends a number of qubits and asks the server to measure them, and the client checks the statistics. Only a server with measurement powers could pass this test, and thus a fully classical machine is eliminated.

Could a purely classical client pass any of the aforementioned tests? Let us begin with Verification 3. A purely classical client can not distinguish between a super fast classical computer and a quantum machine, unless quantum mechanics is wrong. To see this, simply note that whatever

the commands the client sends to the server are, they are classical, and if they represent "quantum" commands then the server can write out the classical description of the entire system and evolve it "on paper" by using the rules of quantum mechanics. The client could never know the difference. Even quantum mechanics being incorrect would not suffice on it's own. For a client to be able to discern between a classical and "quantum" machine, the client would need to know in what sense quantum mechanics is incorrect, and effectively test whether the server (incorrectly) derives measurement outcomes by using a faulty theory.

Concerning Verification 2, this remains an open problem. As we have seen, a crucial technique used in this type of verification is that the client uses *the structure* of the problem at hand to correctly decide the solution with the aid of the server. As we know, if the server is unbounded, using this approach, the client can decide PSPACE problems. This class can certainly simulate all of quantum mechanics and contains BQP. What happens when the server is restricted to BQP will more than likely depend on the structure of BQP problems and how much the client can use that to her advantage.

Since Verification 1 implies both Verification 2 and 3, and Verification 3 cannot be performed by a purely classical client, Verification 1 is impossible for a classical client as well. However, at least two questions remain. Can the flavour of verification appearing in UBQC be relaxed to the point where it verifies something of importance *and* allows for a classical client (an example would be Verification 2, if possible)? In other words, what can we verify with a fully classical client? The other question is, how *much* "quantumness" is required on the side of the client to perform Verification 1? One way of measuring the "quantumness" would be to consider the amount of discord in the individual states the client and server need to share. In UBQC, the states are of the form

$$\sum_{\theta} |\theta\rangle\langle\theta|_{Client} \otimes |+_\theta\rangle\langle+_\theta|_{Server}, \tag{4.1}$$

for the eight angles of θ (plus Z eigenstates needed for trapification). Our initial results show verifiable UBQC can be achieved using only the four BB84 states $|+\rangle, |-\rangle, |0\rangle$ and $|1\rangle$ (this work is in preparation). This could perhaps even be improved to three symmetric states, and still achieve perfect blindness which we need for the verification protocol.

It is clear that blindness and verifiability are intimately linked in UBQC, as blindness enables (but is perhaps not necessary for) verification. As mentioned, our new results (in preparation) suggest that for perfect blindness and verification three symmetric states are optimal (in terms of discord), but *approximate blindness* and verification can be achieved with states of arbitrary small (but non-zero) discord. This, if proven, would constitute a complete characterisation of the necessary resources for Verification type 1, as we have seen that, in this case, a completely classical client (zero discord) is not sufficient.

Next, we focus on the perspective of UBQC as a useful tool in quantum information processing, and other topics which are happening in the UBQC world.

4.2 Other topics

4.2.1 Universal composability

Since UBQC offers private and verifiable *universal* quantum information processing it is not very surprising it could be applied as a "module" in other tasks. Indeed, as we have seen UBQC has already been used in such a way in the proof that $\mathsf{QMIP} = \mathsf{MIP}^*$. Similarly, Mosca and Stebila [82] have proposed UBQC to enable a bank to verify the authenticity of quantum money with a merchant by using only classical communication online (and quantum offline). It is conceivable UBQC could be used for the delegated distribution of highly complicated quantum states to be used as public quantum keys, to be used in other protocols (like Quantum Digital Signatures, see II of this thesis), even when the distributor has only access to a simple device (like single photon emitters, or just a coherent state generator). Many other applications only await discovery. However, the question if UBQC can *at all* be used as a subroutine securely remains an open question, although UBQC itself is *unconditionally secure*. This is the question whether the UBQC protocol is secure in the sense of *universal composability*.

For the details on the topic of the framework of universal composability we refer the reader to papers [83, 84, 85].

The notion of universal composability captures precisely what we wish to be true – that UBQC can safely be used as a subroutine in any other larger protocol, in other words, that it is secure in every context.

The process of proving universal composability of UBQC evolves finding a sensible *composable definition* of security for UBQC, and then prove that the proposed UBQC protocol satisfies the proposed definition. This topic is on-going research, and we do not present all the preliminary results we have in this thesis. However, we note that it seems that the original UBQC is not universally composable in the strictest sense, but with verification added, it may very well be [7]. The verifiability of UBQC is once more emphasized as a key feature.

4.2.2 Alternative models of UBQC

Relatively recently, UBQC has spawned interesting variants with similar functionalities. Here we present a selected few. Recently, Morimae and Fujii [54] have proposed a UBQC variant which uses the Raussendorf, Goyal and Harrington topological model of fault tolerant computation to achieve fault tolerant UBQC. In this work, for the first time, fault tolerance thresholds for UBQC have been calculated for certain standard noise models. Same authors have also proposed three alternative variants of UBQC. In this section we present one of them called "measuring only Alice". In this model, the client has a different power – she is assumed to be capable

[7] During the period between the submission of the First version of this thesis, and the submission of this Final version, a proof idea of composability of the basic UBQC protocol was suggested by Joe Fitzsimons, and is currently being formalized by Portmann, Fitzsimons, Renner, and the author of this thesis.

of performing single qubit measurements only. From a pragmatic point of view, the ability to prepare or measure qubits (photons) are arguably at the same level of difficulty.

In this protocol, the server prepares a large resource state, say the brickwork state, and sends individual qubits to the client, in the sequence of measurements. The client simply measures with respect to the basis which defines the client's desired computation.

This model is conceptually very simple, and one advantage over the original UBQC protocol is that its security is obvious. Since there is no classical communication happening at all, the server learning anything about what the client is doing would violate not only quantum mechanics, but also no–signalling. Also, this protocol is possibly secure in a device–independent setting. There are two disadvantages of the "measuring only Alice" protocol compared to the UBQC scheme. First, photon loss in practice would be a major issue in practice, where it is not an issue at all in the original protocol (Alice can simply resend new photons/qubits until Bob receives them). In "measuring only Alice" this could be resolved by using a thick enough fault-tolerant code, but the losses would in reality be very high, making the scheme impractical. The second disadvantage is that the measurement-only variant of UBQC at this point offers no scheme for verification. However, these issues may be circumvented by modifications of the proposed protocol, and the solutions are in preparation.

On a different note, the relationship between these two protocols, the measurement only and the original UBQC protocol superficially reminds us of the relationship between BB84 and E91 [86, 29] protocols in the history of QKD, which were later proven to be equivalent. While this analogy is not very strong (in the UBQC setting different players send either single qubits, or alternatively, subsystems of an entangled system), it would be interesting to see how far this analogy can be pushed – could we prove the security of the original protocol from the (obvious) security of the measurement-only protocol? A complete equivalence in the sense of BB84 and E91 is unlikely, as the measuring only Alice protocol seems to be secure in the device-independent setting, whereas UBQC does not seem to be.

4.2.3 UBQC and fully homomorphic encryption

As we have noted, the idea of using servers to solve our computationally-intensive problems, while maintaining privacy is by no means restricted to the quantum setting. Perhaps the first formalization of the task as presented in the work *On data banks and privacy homomorhisms* by Rivest, Shamir and Dertouzos, where they posed the problem of finding secure encryption functions, which were also homomorphisms between the sets of plaintexts and cyphertexts with respect to relevant operations – for instance addition and multiplication (for more information please see 1. Very soon after the problem was defined many partial solutions were derived, most often based on various schemes for public-key cryptography. For instance, the RSA, and ElGamal cryptosystems proved to be homomorphisms with respect to multiplication (but not addition). Other schemes allowed the evaluation of both addition and multiplication, but limitations existed in terms of the depth of the computation which can be performed securely [87]. In the game of such secure interactive computation, a very important issue is minimizing rounds of computation

in this model. Namely, the number of communication rounds the client and the server need to perform in order to evaluate a circuit on encrypted data securely.

In 2009. Craig Gentry proposed the first fully homomorphic encryption scheme. Gentry's proposal allows the secure evaluation of any circuit, in such a way that the number of rounds of communication is independent from the size of the circuit. Only the security levels required influence the computational overhead of the encryption and decryption functions, and the security is assured under computational assumptions. The actual problem whose hardness ensures this computational security is the place where Gentry's proposal is somewhat generic. In various proposal which stemmed from his original idea, the security comes from the assumed hardness of the approximate greatest common divisor problem, sparse subset-sum problem, or based on problems on ideal lattices or specific problems over integers [88, 22] [8].

Central to Gentry's proposal is the idea of "bootstrapping" particular partially homomorphic encryptions [22]:

(Informal) *Every (partially) homomorphic encryption scheme E which can evaluate its own decryption circuit, plus a NAND gate can be raised to a (fully) homomorphic scheme E^* which can evaluate any circuit.*

A detailed description how these ideas work is well beyond the scope of this thesis. For this reason, we will refrain from doing so and recommend the interested reader to the accessible expositions of these ideas given by Gentry in [89]. Here, we shall give the "big picture" comparison of UBQC and Gentry's result.

Global functionality: In FHE, what is concealed is the input and the generated output of the computation, whereas the function which is to be evaluated is public. In UBQC the function is hidden as well (but can also be public). However, using a construction called "Yao's garbled circuits" [22] the function can be concealed in FHE as well. Larger differences become obvious once the allowed requirements on the client are analysed. This we address presently.

Security levels: The security in UBQC is information-theoretic, assuming the validity of quantum mechanics. The security in FHE is by construction based on a public-key cryptosystem and all classical public-key cryptosystems are secure under computational assumptions by construction [9]. However, there exist schemes which are unconditionally secure (under cyphertext-only attacks) partial homomorphic encryptions, and it is unclear whether FHE could be extended into a symmetric-key unconditionally secure scheme.

Rounds of communication: As we have noted, the number of communication rounds in FHE is independent of the size of the circuit to be evaluated. On the other hand, in UBQC a commu-

[8] It is true, however, that some of these schemes were ultimately broken – proven not to be secure.

[9] More precisely, what they achieve in FHE is *semantic security* which is a particular type of guarantee defined for asymmetric key cryptosystems, equivalent to ciphertext indistinguishability. Handwaveingly, the latter implies that an adversary, having access to a public key, for any two plaintexts of his own choosing, having received a cyphertext which is an encryption of one of the plaintexts cannot guess which of the two plaintexts the cyphertext encrypts with probability (significantly) larger than $1/2$. This type of security is the most basic of requirements in public key schemes, as the adversary has the power to encrypt texts of his own choosing, and thus generate pairs $(plaintext, cyphertext)$. This in turn implies that such encryption schemes have to be non-deterministic.

nication exchange is needed for each measurement, equivalently, for each gate in the quantum circuit which is to be evaluated. In this sense, the differences between two schemes are vast, and the reason why they both nonetheless exciting from the practical point of view is because they ultimately solve different and yet related problems.

In the setting of UBQC, the problem we are addressing is as follows: can a client, with access to simple devices and a classical computer, evaluate quantum functions with the help of the server in such a way that this delegated process is a fair substitute to the client having a quantum computer of her own. The crucial difference between the client and the server is in the type of functions they are capable evaluating on their own in a very broad sense. The client can evaluate classical and the server quantum functions, and the actual number of operations needed to do so is then simply bounded to be polynomial in the instance size. One of the key assumptions here is that everything "classical" (classical computation and communication) or "almost classical" (single qubit generation) is almost free compared to the difficulties in coherently manipulating many qubits. This assumption is, at least to scientist who are trying to find ways to build a quantum computer, obvious and justified.

In the setting of FHE, the client and the server are, in the broad sense addressed in UBQC, computationally equivalent. Both are capable of evaluating the same types of functions, it is just the case the server can do it a lot more rapidly (or perhaps can handle a lager instance of the problem) then the client in terms of *real elapsed time*. Here, it is then crucial that the number of rounds of communication be limited, compared to the actual computation size.

From a purely practical point of view, the client's powers in UBQC are already modest. Moreover, following the results we presented in this work, namely that the required security levels can be guaranteed even in the presence of realistic imperfections, only limited practical advantages can be made by making the client more, or purely classical [10]. In the modern age, the classical communication is mostly done over optical cables, and the already existing mediums and technology could be used to distribute single photons (or weak coherent pulses) from the client to the server. An important advancement would be made if the number of communication and processing rounds used in UBQC could be made relatively small compared to the computation size. This we place as one of the main challenges in the future of UBQC. If such a goal could be achieved, then UBQC could indeed be compared to FHE on level grounds, with the clear advantage of UBQC being capable of resolving BQP problems. Otherwise, we maintain that UBQC and FHE should not be directly compared due to the fundamental differences in their designed functionality.

4.3 Reproving blindness, and the relationship between UBQC and MBQC

In this section we present a more formal proof of blindness. Additionally we analyse which specific properties of the one-way model are responsible for the blindness of UBQC.

[10] This question is, on the other hand, one of the most exciting questions from a theoretical point of view, as we have argued earlier in this chapter.

The definition of blindness states that the system of the server's register, at any stage of the computation should be independent from the classical and quantum information he receives from the computational angles (ϕ_i). A consequence of this is that even in the presence of server's prior knowledge about the computational angles the same claim holds. In particular, all proofs of blindness have to work, independent of the existence of prior knowledge. However, in all the proofs presented it was assumed that the server cannot learn anything about the client's hidden r parameters: in the original proof it was shown that the state the server receives at the i^{th} step of the protocol is a maximally mixed state, *once the clients secret bits* r have been traced out. Similarly, in other proofs, we assumed that the server's responses cannot influence the client's behaviour, in which case we can pre-set the servers responses to be fixed (and independent from the states he receives from the client), and observe the entirety of information the server gets throughout the protocol, after which, by tracing out the r parameters we again obtain a maximally mixed state in the hands of the server. This assumption again relies on the server never obtaining any information about the hidden r parameters. And it is, strictly speaking, false.

To see this, simply consider the setting where the server knows exactly the first measurement angle ϕ_1 and nothing else. Then, he receives the angle $\delta_1 = \phi_1 + r_1\pi + \theta_1$ and is in the possession of a qubit in the state $|+_\theta\rangle$. Clearly, by measurement of the qubit in the basis $|\pm_{\delta-\theta}\rangle$, the outcome of the measurement is exactly r_1. Thus, a general proof of blindness should not assume all the r parameters can be traced out from the server's system [11] What we present now is a more detailed proof of blindness [12]. Fortunately, this proof shows not only that UBQC satisfies blindness, but also that all the presented results are *morally* correct, that is, the hidden assumptions can be justified, and have thus been presented in this thesis in their original form.

4.3.1 Proof of blindness

For simplicity of notation we will often be omitting the symbol for the tensor product $\otimes$, the pure state density operators will be represented by boxes *i.e.* $|\phi\rangle\langle\phi| = \boxed{\phi}$ and we will use a one dimensional indexing for qubits, angles and binary variables, ranging from 1 to N where N is the computation size. We begin by presenting the proof for the setting with no prior knowledge about the computational angles. With $\mathbb{1}$ we denote the identity operator, and assume its dimension is clear from context. Subsystems are often swapped for convenience reasons. We explicitly write out the states of Bob's system at particular stages of the UBQC protocol:

[11] This subtle problem was pointed out by Renato Renner and Matthias Christandl during a visit to ETH Zurich in June 2012., and ultimately resolved there.

[12] His proof was derived from the discussions with Fernando Brandao, Christopher Portmann and Anthony Leverrier. To this group I extend my profound gratitude for their involvement, interest, time and patience in the problems addressed in this section.

$$
\begin{aligned}
(post-preparation) \qquad & \sum_{\theta_1,\ldots,\theta_N} \bigotimes_{i=1}^{N} \boxed{+_{\theta_i}} && (= \mathbb{1}/2^N) \\
(1\ Alice's\ round) \qquad & \sum_{r_1=0}^{1} \sum_{\theta_1,\ldots,\theta_N} \bigotimes_{i=1}^{N} \boxed{+_{\theta_i}} \otimes \boxed{\delta 1} && = \\
& \sum_{r_1=0}^{1} \sum_{\theta_1} \boxed{\delta_1} \otimes \boxed{+_{\theta_1}} \bigotimes_{i=2}^{N} \sum_{\theta_i} \boxed{+_{\theta_i}} && = \\
& \sum_{r_1=0}^{1} \sum_{\theta_1} \boxed{\delta_1} \otimes \boxed{+_{\theta_1}} \otimes \mathbb{1}/2^{N-1} && = \\
& \sum_{r_1=0}^{1} \sum_{\theta_1} \boxed{\phi_1 + r_1\pi + \theta_1} \otimes \boxed{+_{\theta_1}} \otimes \mathbb{1}/2^{N-1} && =
\end{aligned}
\tag{4.2}
$$

Now we show that $\sum_{r_1=0}^{1} \sum_{\theta_1} \boxed{\phi_1 + r_1\pi + \theta_1} \otimes \boxed{+_{\theta_1}}$ is the totally mixed state. We define the following four-qubit unitary transformation CCU (standing for "classically controlled unitary"):

$$
CCU \boxed{\varphi} \otimes |\psi\rangle\langle\psi| CCU^\dagger = \boxed{\varphi} \otimes Z_{-\varphi}|\psi\rangle\langle\psi|Z_{-\varphi}^\dagger \tag{4.3}
$$

where φ is one of the eight angles appearing in the UBQC protocol, encoded into a three qubit state. Now consider the following state:

$$
\begin{aligned}
& CCU \sum_{r_1=0}^{1} \sum_{\theta_1} \boxed{\phi_1 + r_1\pi + \theta_1} \otimes \boxed{+_{\theta_1}} CCU^\dagger = \\
& \sum_{r_1=0}^{1} \sum_{\theta_1} CCU \boxed{\phi_1 + r_1\pi + \theta_1} \otimes \boxed{+_{\theta_1}} CCU^\dagger = \\
& \sum_{r_1=0}^{1} \sum_{\theta_1} \boxed{\phi_1 + r_1\pi + \theta_1} \otimes \boxed{+_{\theta_1 - (\phi_1 + r_1\pi + \theta_1)}} = \\
& \sum_{r_1=0}^{1} \sum_{\theta_1} \boxed{\phi_1 + r_1\pi + \theta_1} \otimes \boxed{+_{-\phi_1 - r_1\pi}} = \\
& \sum_{r_1=0}^{1} \boxed{+_{-\phi_1 - r_1\pi}} \otimes \left(\sum_{\theta_1} \boxed{\phi_1 + r_1\pi + \theta_1} \right) = \\
& \sum_{r_1=0}^{1} \boxed{+_{-\phi_1 - r_1\pi}} \otimes \mathbb{1}/2^3 = \mathbb{1}/2^4
\end{aligned}
\tag{4.4}
$$

Since CCU is unitary $\sum_{r_1=0}^{1} \sum_{\theta_1} \boxed{\phi_1 + r_1\pi + \theta_1} \otimes \boxed{+_{\theta_1}}$ is a totally mixed state.

So we continue with the derivation of Bob's state after Alice's first round:

$$\sum_{r_1=0}^{1}\sum_{\theta_1}\boxed{\phi_1+r_1\pi+\theta_1}\otimes\boxed{+_{\theta_1}}\otimes\mathbb{1}/2^{N-1} = \mathbb{1}/2^4\otimes\mathbb{1}/2^{N-1} = proport.\ to\ \mathbb{1} \tag{4.5}$$

From here on out we will be omitting the normalization factors. Next, we consider Bob's state after his first response allowing that he applied an arbitrary instrument on his system. For completeness we write out all the previous stages:

$$\begin{aligned}
(post-preparation) \quad & \sum_{\theta_1,\ldots,\theta_N}\bigotimes_{i=1}^{N}\boxed{+_{\theta_i}} = \mathbb{1}/2^N \\
(1\ Alice's\ round) \quad & \sum_{r_1=0}^{1}\sum_{\theta_1,\ldots,\theta_N}\bigotimes_{i=1}^{N}\boxed{+_{\theta_i}}\otimes\boxed{\delta_1} = \mathbb{1}/2^N \\
(1\ Bob's\ round) \quad & \sum_{s_1^B=0}^{1}\sum_{r_1=0}^{1}\sum_{\theta_1,\ldots,\theta_N}\boxed{s_1^B}\otimes\mathcal{E}^{s_1^B}\left(\bigotimes_{i=1}^{N}\boxed{+_{\theta_i}}\otimes\boxed{\delta_1}\right) = \\
& \sum_{s_1^B=0}^{1}\boxed{s_1^B}\otimes\mathcal{E}^{s_1^B}(\mathbb{1}) = \\
& \sum_{s_1^B=0}^{1}\mathcal{E}'^{s_1^B}\left(\boxed{s_1^B}\otimes\mathbb{1}\right) = \\
& \sum_{s_1^B=0}^{1}\mathcal{E}''^{s_1^B}(\mathbb{1}) = \mathcal{E}'''(\mathbb{1})
\end{aligned}$$

We now skip (in writing out) Alice's second round and proceed directly to Bob's second round, starting from the explicit state in (4.6)

$$(1\ Bob's\ round) \quad \sum_{s_1^B=0}^{1}\sum_{r_1=0}^{1}\sum_{\theta_1,\ldots,\theta_N}\boxed{s_1^B}\otimes\mathcal{E}^{s_1^B}\left(\bigotimes_{i=1}^{N}\boxed{+_{\theta_i}}\otimes\boxed{\delta_1}\right)$$

$$(2\ Bob's\ round)$$

$$\begin{aligned}
& \sum_{s_1^B,s_2^B}\sum_{r_1,r_2}\sum_{\theta_1,\ldots,\theta_N}\boxed{s_2^B}\otimes\mathcal{E}^{s_2^B}\left(\boxed{s_1^B}\otimes\boxed{\delta_2}\otimes\mathcal{E}^{s_1^B}\left(\bigotimes_{i=1}^{N}\boxed{+_{\theta_i}}\otimes\boxed{\delta_1}\right)\right) = \\
& \sum_{s_1^B,s_2^B}\sum_{r_1,r_2}\sum_{\theta_1,\ldots,\theta_N}\boxed{s_2^B}\otimes\mathcal{E}^{s_2^B}\left(\mathcal{E}'^{s_1^B}\left(\boxed{s_1^B}\otimes\boxed{\delta_2}\otimes\bigotimes_{i=1}^{N}\boxed{+_{\theta_i}}\otimes\boxed{\delta_1}\right)\right) =
\end{aligned} \tag{4.6}$$

$$\sum_{s_1^B,s_2^B}\sum_{r_1,r_2}\sum_{\theta_1,\ldots,\theta_N}\mathcal{E}^{s_1^B,s_2^B}\left(\boxed{s_2^B}\otimes\boxed{s_1^B}\otimes\boxed{\delta_2}\otimes\bigotimes_{i=1}^{N}\boxed{+_{\theta_i}}\otimes\boxed{\delta_1}\right)=$$

$$\sum_{s_1^B,s_2^B}\mathcal{E}^{s_1^B,s_2^B}\left(\boxed{s_2^B}\otimes\boxed{s_1^B}\otimes\sum_{r_1,r_2}\sum_{\theta_1,\ldots,\theta_N}\left(\bigotimes_{i=1}^{N}\boxed{+_{\theta_i}}\otimes\boxed{\delta_1}\otimes\boxed{\delta_2}\right)\right)=$$

$$\sum_{s_1^B,s_2^B}\mathcal{E}^{s_1^B,s_2^B}\left(\boxed{s_2^B}\otimes\boxed{s_1^B}\otimes\sum_{r_1,r_2}\sum_{\theta_1,\theta_2}\left(\boxed{+_{\theta_1}}\otimes\boxed{\delta_1}\otimes\boxed{+_{\theta_2}}\otimes\boxed{\delta_2}\right)\otimes\mathbb{1}\right)=$$

Again, one now shows that $\sum_{r_1,r_2}\sum_{\theta_1,\theta_2}\left(\boxed{+_{\theta_1}}\otimes\boxed{\delta_1}\otimes\boxed{+_{\theta_2}}\otimes\boxed{\delta_2}\right)=\mathbb{1}$. We present a more general proof of this claim, which we will use in the k^{th} for the analysis of the k^{th} step of the protocol:

$$\sum_{r_1,\ldots,r_k}\sum_{\theta_1,\ldots,\theta_k}\left(\boxed{+_{\theta_1}}\otimes\boxed{\delta_1}\otimes\ldots\otimes\boxed{+_{\theta_k}}\otimes\boxed{\delta_k}\right)$$

Note that r_l and θ_l do not appear as parameters in δ_j for $j<l$. Thus,

$$\sum_{r_1,\ldots,r_{k-1}}\sum_{\theta_1,\ldots,\theta_k-1}\left(\boxed{+_{\theta_1}}\otimes\boxed{\delta_1}\otimes\ldots\otimes\boxed{+_{\theta_{k-1}}}\otimes\sum_{r_k}\sum_{\theta_k}\left(\boxed{+_{\theta_k}}\otimes\boxed{\delta_k}\right)\right)$$

Recall $\delta_k=\phi'_k+\theta_k+r_k\pi$. The variable ϕ'_k depends on the values s^B, and so does δ, so the reporting activity of the server is taken into account. However, ϕ'_k does not depend on r_k (or θ_k) so again by the "CCU" argument used in (4.4) we have that $\sum_{r_k}\sum_{\theta_k}\left(\boxed{+_{\theta_k}}\otimes\boxed{\delta_k}\right)=\mathbb{1}$. Since this holds for all k, by iteration we get:

$$\sum_{r_1,\ldots,r_k}\sum_{\theta_1,\ldots,\theta_k}\left(\boxed{+_{\theta_1}}\otimes\boxed{\delta_1}\otimes\ldots\otimes\boxed{+_{\theta_k}}\otimes\boxed{\delta_k}\right)=\mathbb{1}\tag{4.7}$$

To finish Bob's second round, we then have:

$$(2\ Bob's\ round)$$

$$\sum_{s_1^B,s_2^B}\sum_{r_1,r_2}\sum_{\theta_1,\ldots,\theta_N}\boxed{s_2^B}\otimes\mathcal{E}^{s_2^B}\left(\boxed{s_1^B}\otimes\boxed{\delta_2}\otimes\mathcal{E}^{s_1^B}\left(\bigotimes_{i=1}^{N}\boxed{+_{\theta_i}}\otimes\boxed{\delta_1}\right)\right)=$$

$$\sum_{s_1^B,s_2^B}\mathcal{E}^{s_1^B,s_2^B}\left(\boxed{s_2^B}\otimes\boxed{s_1^B}\otimes\sum_{r_1,r_2}\sum_{\theta_1,\theta_2}\left(\boxed{+_{\theta_1}}\otimes\boxed{\delta_1}\otimes\boxed{+_{\theta_2}}\otimes\boxed{\delta_2}\right)\otimes\mathbb{1}\right)=$$

$$\sum_{s_1^B,s_2^B}\mathcal{E}^{s_1^B,s_2^B}\left(\boxed{s_2^B}\otimes\boxed{s_1^B}\otimes\mathbb{1}\otimes\mathbb{1}\right)=\mathcal{E}(\mathbb{1})$$

This approach is generalized for the k^{th} step which we explicitly write out:

$$
\begin{aligned}
&(k\ Bob's\ round) \\
&\sum_{r_1,\ldots,r_k}\sum_{\theta_1,\ldots,\theta_k}\sum_{s_1^B,\ldots,s_k^B} \boxed{s_k^B} \otimes \\
&\mathcal{E}^{s_k^B}\left(\boxed{\delta_k} \otimes \boxed{s_{k-1}^B} \otimes \mathcal{E}^{s_{k-1}^B}\left(\boxed{\delta_{k-2}} \otimes \right.\right. \\
&\left.\left.\boxed{s_{k-2}^B} \cdots \boxed{s_1^B} \otimes \mathcal{E}^{s_1^B}\left(\boxed{\delta_1} \otimes \boxed{+_{\theta_1}} \otimes \cdots \otimes \boxed{+_{\theta_k}}\right)\cdots\right)\right)
\end{aligned}
\tag{4.8}
$$

We now take out all the CP maps (by modifying them to act trivially on the systems they had not been acting before)and collect them in one single CP map:

$$
\begin{aligned}
&\sum_{r_1,\ldots,r_k}\sum_{\theta_1,\ldots,\theta_k}\sum_{s_1^B,\ldots,s_k^B} \mathcal{E}^{s_1^B,\ldots,s_k^B}\left(\bigotimes_{i=1}^{k}\boxed{s_i^B} \otimes \boxed{\delta_i} \otimes \boxed{+_{\theta_i}}\right) = \\
&\sum_{s_1^B,\ldots,s_k^B} \mathcal{E}^{s_1^B,\ldots,s_k^B}\left(\bigotimes_{i=1}^{k}\boxed{s_i^B} \otimes \sum_{r_1,\ldots,r_k}\sum_{\theta_1,\ldots,\theta_k}\left(\bigotimes_{i=1}^{k}\boxed{\delta_i} \otimes \boxed{+_{\theta_i}}\right)\right)
\end{aligned}
$$

By the equation (4.7) we then have:

$$
(k\ Bob's\ round)
$$

$$
\sum_{s_1^B,\ldots,s_k^B} \mathcal{E}^{s_1^B,\ldots,s_k^B}\left(\bigotimes_{i=1}^{k}\boxed{s_i^B} \otimes \mathbb{1}\right) = \sum_{s_1^B,\ldots,s_k^B} \mathcal{E}'^{s_1^B,\ldots,s_k^B}(\mathbb{1}) = \mathcal{E}(\mathbb{1}) \tag{4.9}
$$

Thus, Bob's system is always independent from the computational angles $\{\phi_i\}$, stated as the following theorem:

Theorem 6. *In the UBQC protocol with a classical input the state of Bob's system is independent from Alice's computational angles at each step of the protocol.*

The statement of the theorem above could be used as a definition of blindness as well, specialized to the MBQC computation case. The proof could most likely easily be extended to handle quantum inputs as well. Note that the analysis above also shows that Bob's response strategy is inconsequential for security. As we have shown, Bob's system, regardless of his actions is always independent from the choice of Alice's angles. Thus, in the case of prior information his information cannot increase through the run of the protocol.

4.3.1.1 Approximate blindness

In Chapter 2 we considered the case of qubit preparation errors. Of particular relevance was the case the emitted states generated by the client are general states ρ^θ instead of $|+_\theta\rangle$, which assumes the quantum state generator has no memory and the angles are chosen by a perfect source of randomness. With these assumptions we have shown that the preparation error directly

corresponds to overall security levels, as show in inequality 10.77. Here, we derive the same result by consider the state of Bob's system at the k^{th} step in the ideal and non ideal case. We will denote the ideal state (for the k^{th} step) π^{ideal}, given in equation below 4.8, and the realistic state for the same state is then given with:

$$\pi_{real} = \sum_{r_1,\ldots,r_k} \sum_{\theta_1,\ldots,\theta_k} \sum_{s_1^B,\ldots,s_k^B} \boxed{s_k^B} \otimes \mathcal{E}^{s_k^B}\left(\boxed{\delta_k} \otimes \boxed{s_{k-1}^B} \otimes \mathcal{E}^{s_{k-1}^B}\left(\boxed{\delta_{k-2}} \otimes \boxed{s_{k-2}^B} \cdots \boxed{s_1^B} \otimes \mathcal{E}^{s_1^B}\left(\boxed{\delta_1} \otimes \rho^{\theta_1} \otimes \cdots \otimes \rho^{\theta_k}\right)\cdots\right)\right)$$

By the linearity of the trace norm (also a few triangle inequalities and contractivity of CPTP maps – $\|\mathcal{E}\rho - \mathcal{E}(\nu)\|_{tr} \leq \|\rho - \nu\|_{tr}$), it is easy to see that

$$\frac{1}{2}\|\pi_{ideal} - \pi_{real}\|_{tr} \leq N\epsilon_{prep}, \tag{4.10}$$

where $\epsilon_{prep} = \max_\theta \frac{1}{2}\||+_\theta\rangle\langle+_\theta| - \rho^\theta\|_{tr}$.

4.3.2 Properties of the one-way model and UBQC

The one-way computation model over the family of brickwork states, or any other family of universal resource states, with measurements in the XY plane of the Bloch sphere is universal. Once the family of resource states has been fixed, the degree of freedom which defines the computation are the measurement angles, and they are the key information which needs to be hidden from the server. Here, we will show that the description of computation in this model, *i.e.* the measurement angles, allow for a large amount of redundancy. By redundancy here we mean that many differing sets of computational angles define the same computation.

Such redundancy appears in he circuit model, and stems from the fact that a particular $n-$qubit unitary can be decomposed into single and two qubit gates in many ways. The redundancy we will show in the one-way model stems from the need to counteract the probabilistic nature of measurement in order to achieve universality. This type of redundancy is specific to measurement-based models of quantum computation, and we will argue that it is a crucial component in the blindness properties of UBQC.

First, recall the one-way map $\Gamma_G(\phi_{kk})$ which we defined as follows:

Definition 16. *Let G be an open graph state, over N vertices with n and m vertices in the input and output partitions, and let $(\phi_1, \ldots, \phi_{N-m})$ be a sequence of computational angles. Then the one-way map*

$$\Gamma_G(\phi_1, \ldots, \phi_{N-m}) \tag{4.11}$$

is the linear map mapping the Hilbert space of the input partition to the output partition in a

run of a postselected one-way computation, where all the measurements collapse to the positive branch.

The one-way computation we refer in the definition above was defined by Protocol 1.2.1. A key property of the one-way model is that, provided the entanglement geometry of the underlying graph state has particular properties as flow or gflow, then the map Γ can be implemented deterministically by adapting the measurement angles [35, 39]. This has been illustrated in Protocol 3.

In the illustrations of corrective strategies for the one-way model in Section 1.2.2 we have shown that the global state of the one-way computer in the two following scenarios are equal:

- after measurement of a qubit v with measurement angle ϕ_v, with the (non-positive branch) outcome $s_v = 1$
- after measurement of a qubit v with measurement angle $\phi_v + \pi$, with the (positive branch) outcome $s_v = 0$.

Since the desired overall one-way map can be insured by adaptive measurement is the first scenario above, the same adaptive measurement strategy can help regain the desired one-way map if we perform a measurement angle shifted by π.

The consequence of this is redundancy as shown by the following lemma:

Lemma 17. *Let G be an open graph state with flow over N qubits, let m be the size of the output partition and let l be the number of qubits in the last and second to last non-output layers. Let $\overrightarrow{r} = (r_1, \ldots, r_{N-m}) \in \{0,1\}^{\times N-m-l}$ be a sequence of bits, such that $r_{N-m-l+1}, \ldots, r_{N-m} = 0$, so that the r values corresponding to the labels of the last and second to last non-output layers are always set to zero. Then for every sequence of measurement angles $(\phi_1, \ldots, \phi_{N-m})$ the following holds:*

$$\Gamma_G(\phi_1, \ldots, \phi_{N-m}) = \Gamma_G(\phi'_1 + r_1\pi, \ldots, \phi'_{N-m} + r_{N-m}\pi),\ \forall \overrightarrow{r}, \tag{4.12}$$

if

$$\phi'_k = (-1)^{\oplus_{j \in D_k^X} r_j} \phi_k + (\bigoplus_{j \in D_k^Z} r_j)\pi. \tag{4.13}$$

The dependency sets D_k^X and D_k^Z are given by the flow construction as explained earlier.

Note that the angles ϕ'_i depend on $\overrightarrow{r}$ so for reasons of clarity we will write $\phi'_i(\overrightarrow{r})$. However the angle ϕ'_i does not depend on the digit r_i for any i, as it only depends on r_j where $j < i$, by the flow construction. Thus he sequences $\{\phi'_i(\overrightarrow{r}) + r_i\pi\}$ and $\{\phi'_i(\overrightarrow{r}') + r'_i\pi\}$ are distinct for $\overrightarrow{r} \neq \overrightarrow{r'}$. This can be shown by induction. This means that, due to the structure of the one-way model itself, there are $2^{(N-m-l)}$ distinct sequences of measurement angles which realize the same computation. Note that this redundancy *necessarily* has to exist to enable the adaptive structure to combat undesired measurement outcomes and is intrinsic to the one-way model itself.

In Lemma 17 the requirement $r_{N-m-l+1} = \cdots = r_{N-m} = 0$ ensured the realized one-way maps were identical. However, if we relax this constraint and allow these latter r values to attain arbitrary values, the realized map attains the form:

$$\left(\bigotimes_{v=N-m}^{N} X_v^{\oplus_{j\in D_v^X} r_j} Z_v^{\oplus_{j\in D_v^Z} r_j} \right) \Gamma_G(\phi'_1 + r_1\pi, \ldots, \phi'_{N-m} + r_{N-m}\pi) \tag{4.14}$$

That is, we obtain the original one-way map with an additional operator applied on each output qubit, which depends on the values of r_i pertaining to the labels of the last and second to last non-output layer. To see this, please refer to last step of Protocol 3.

In the setting of UBQC all the computational angles come from the set $\phi \in \left\{ \frac{k\pi}{4} \right\}_{k=0}^{7}$. Then, each angle ϕ can be encoded as three binary digits q^1, q^2 and q^3 as

$$\phi = q^1\pi + q^2\frac{\pi}{2} + q^3\frac{\pi}{4}. \tag{4.15}$$

Then the observation given in equation 4.14 has the following consequence:

Lemma 18. *Let G be the brickwork state, with the leftmost column as the input layer, and rightmost as the output, and let $(\phi_1, \ldots, \phi_{N-n})$ be a sequence of measurement angles from the discrete set $\left\{ \frac{k\pi}{4} \right\}_{k=0}^{7}$. Let (q^1, q^2, q^2) be a binary representation of the angles in the discrete set (as described above), and let*

$$\mathcal{C}_{(\phi_1,\ldots,\phi_{N-n})} = \left\{ \left(\bigotimes_{v=N-m}^{N} X_v^{x_v} Z_v^{z_v} \right) \Gamma_G(\phi_1, \ldots, \phi_{N-n}) | x_v, z_v \in \{0,1\}, \forall v \right\} \tag{4.16}$$

be the set of all (unitary) maps equal to $\Gamma_G(\phi_1, \ldots, \phi_{N-n})$ up to local Pauli operators. Then for any sequence of most significant digits $\{q_1^1, \ldots q_{N-n}^1\}$ there exists a sequence of measurement angles $(\phi'_1, \ldots, \phi'_{N-n})$ such that

1. *The most significant digit of ϕ'_k is q_k^1 for all k, and*
2. *$\Gamma_G(\phi'_1, \ldots, \phi'_{N-n}) \in \mathcal{C}$.*

This means that if we are interested in implementing a unitary over a brickwork state only up to local Pauli corrections then the most significant digit of the computational measurement angles we need to perform is *independent from the intended computation*.

Again this is a consequence of the intrinsic redundancy of the one-way model, and incidentally could be understood from the arbitrariness of which sequence of measurement outcomes we *choose* to use as the computation-defining one, as we have already briefly mentioned. By convention only this is the positive branch (which assumes the reference frame for the qubit states is fixed in a particular way).

With this observation we can explain blindness in a slightly different way. In UBQC, the client wishes to implement some unitary on the server, however to maintain the privacy of the output, the output has to be quantum one-time padded. This will also induce a (classical) one-time pad

on classical measurement outcomes, should the client want a classical outcome.

However, a quantum one-time pad is just a process in which random local Pauli operators are applied to a quantum system. So, if a client actual desired computation is $\Gamma_G(\phi_1, \ldots, \phi_{N-n})$, to maintain the privacy, the client has to implement a random computation in the set $\mathcal{C}$ defined in Lemma 18. This random computation in $\mathcal{C}$ can be again written in the form of a one-way map for some sequence of measurement angles: $\Gamma_G(\phi'_1, \ldots, \phi'_{N-n})$, however, by virtue of Lemma 18, now the most significant digits of the implemented angles ϕ'_k are independent from the clients actual desired computation.

Next, we consider which computation the server performs in a particular run of a UBQC protocol, and what information he has access to. The systems which the server obtains from the client for each measurement in UBQC are the state $|+_\theta\rangle$ and the angle $\delta = \phi' + \theta + r\pi$ (for simplicity we are omitting the indices). This pair of classical and quantum information is up to a classically controlled unitary equivalent to the state $|+_{\phi'+r\pi}\rangle$, which is correlated to the actual computational angle, and $\delta = \phi' + \theta + r\pi$, which no longer is. This holds because nothing in the server's system is correlated to θ and thus δ is no longer correlated to the computational angle, even given the quantum state $|+_{\phi'+r\pi}\rangle$. This action of a local unitary preserves the correlations between the client's and server's systems.

The value ϕ' depends on the server's prior responses, and the server knows these, so we may pre-set them to zero for simplicity. The angles $\phi' + r\pi$ have the property that they implement the desired unitary of the client up to local Pauli corrections, that is a map in the set $\mathcal{C}_{(\phi_1,\ldots,\phi_{N-n})}$ as defined in Lemma 18 for the client's actual computational angles $\phi_1, \ldots, \phi_{N-n}$. By the same lemma the most significant digit of the angles $\phi' + r\pi$ is independent from the actual computation desired by the client, but the less significant digits define the computation (up to local Pauli corrections).

As we have shown the server has "one-qubit" worth of information about the computation the client wishes, contained in the state $|+_{\phi'+r\pi}\rangle$. However, the server can only access the information about the most significant digit about the angle $\phi' + r\pi$ from the state $|+_{\phi'+r\pi}\rangle$, which is independent from the computation itself. The less significant digits are inaccessible. To see this consider a state $|+_\theta\rangle$ where θ is chosen uniformly at random form the discrete set of eight "UBQC states". The angle θ can then be represented in the binary representation (q^1, q^2, q^3) and the state itself written as $|+_\theta\rangle\langle+_\theta| = Z^{q^3}_{\pi/4} Z^{q^2}_{\pi/2} Z^{q^1} |+\rangle\langle+| Z^{q^1\dagger} Z^{q^2\dagger}_{\pi/2} Z^{q^3\dagger}_{\pi/4}$.

Assume a POVM, characterised by the elements $\{\Pi_j\}_j$ is performed on this qubit, and we are interested in the a-posteriori distribution of q^2 and q^3, which is initially uniform. Assume some outcome k corresponding to the element Π_k has occurred. We have the following derivation:

$$P(q^2, q^3|k) = \frac{P(k|q^2, q^3)P(q^2, q^2))}{P(k)} = \frac{Tr(\Pi_k \sum_{q_1} Z^{q^1} \left(Z^{q^3}_{\pi/4} Z^{q^2}_{\pi/2} |+\rangle\langle+| Z^{q^2\dagger}_{\pi/2} Z^{q^3\dagger}_{\pi/4} \right) Z^{q^1\dagger}) P(q^2, q^2)}{Tr(\Pi_k)} \tag{4.17}$$

Note that for any angle ω it holds that $\sum_{q_1} Z^{q^1}|+_\omega\rangle\langle+_\omega|Z^{q^1\dagger} = \mathbb{1}$, so we have

$$P(q^2, q^3|k) = \frac{Tr(\Pi_k)P(q^2, q^2)}{Tr(\Pi_k)} = P(q^2, q^3). \tag{4.18}$$

Colloquially, this means that, provided no prior knowledge about the most significant digit of angle θ is available, no measurement of the quantum state $|+_\theta\rangle$ can increase our knowledge about the less significant binary digits of θ. However, as we have seen, the client has complete freedom to choose the most significant digit uniformly at random.

The analysis above is not meant to produce yet another proof of blindness nor should it be taken as such. It is just a semi-formal discussion potentially clarifying one perspective which links blindness properties of UBQC and the intrinsic properties of the one-way model.

To summarize, we have attempted to argue that the redundancy of the one-way model and information-theoretical limits of the powers of measurements are a key component of the blindness properties of UBQC. The probabilistic nature of the one-way model due which the redundancy we discussed in this section exists also gives rise to the quantum one-time paddedness of the output layer. But, probabilistic measurement outcomes are unavoidable in MBQC in general – to insure perfect flow of information from the measured qubit to the rest of the system, the entanglement has to be maximal. But then, the reduced density operator for each qubit prior to measurement has to be totally mixed. While being a difficulty with respect to the correctness of computation, which has to be adaptive (thereby, for instance, disallowing single-shot measurement instantaneous computations – depth 1 universal computations), probabilistic measurement outcomes in MBQC allow for interesting cryptographic applications like UBQC. This is certainly one of the reasons why the notion of UBQC was much more likely to be developed in this model of quantum computation rather than in the circuit model or other alternatives.

Chapter 5

Side results: *generalized phase map decomposition*

We analyse the structure of the matrix of the linear map, implemented in the positive branch of a one-way computation. From this analysis certain known results on when a quantum computation can be simulated in the strong sense are recaptured using different techniques.

In universal blind quantum computation, the secret quantum program the client conceals from the server is characterised by the underlying graph defining the structure of the computational resource state, and the measurement angles of the positive branch.

The computational angles could have been generated by the client Alice from a quantum circuit by using one of many translation techniques between the models of MBQC and quantum circuits. However, in principle, it is possible to "program" a measurement based quantum computer by directly thinking in terms of the measurement angles, without going through the circuit model. For this to be possible one has to understand the relationship between the geometry of the resource graph state and the computation angles on one side, and the actual map which they realize in the process of computation. Conversely, one may attempt to construct a measurement-based computation which realizes a given unitary map. In this chapter we study the relationship between the description of a measurement-based quantum computation and the map it implements when run, in terms of the matrix representing the map in the computational basis.

This question is not relevant only from the perspective of UBQC, but also for MBQC itself. The remainder of this chapter is thus presented in the more general context independent from UBQC, and can be read independently from the rest of the thesis.

5.1 Introduction

The one-way model of quantum computation has drawn considerable attention, mainly because it suggests different physical realisations of quantum computing [24, 25]. Recall, in this model quantum states are transformed using single qubit measurements on an entangled state (called *open graph state*), which is prepared from an input state by performing controlled-Z operations on pairs of qubits, including the input system and auxiliary qubits prepared in the $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ state. Quantum measurements are probabilistic in general, and can drive the computation over 2^n different branches, where n is the number of measurements. However,

there exist sufficient conditions based on the structure of the graph state where the computation can be controlled by means of single qubit corrections, dependent on the previous measurement outcomes, so that the entire computation becomes deterministic [24, 35, 39, 25]. In such a deterministic computation, all the branches implement the same unitary map introduced by the *positive branch* (also known as the post-selected branch) which corresponds to the scenario in which every measurement collapses the qubit states to pre-selected states, typically $|+_{\alpha_i}\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\alpha_i}|1\rangle)$.

We give a complete structural characterisation of the map the positive branch of a one-way pattern implements. The positive branch of a one-way pattern can be expressed in terms of a *phase map decomposition* $R\Phi P$ [90, 91], which we then further analyse to obtain the primary structure of the matrix M which represents $R\Phi P$ in the computational basis. The columns of M can be written as:

$$M\mathbf{e}_i = \varepsilon_i B_i \overrightarrow{\varphi} \tag{5.1}$$

where ε_i are complex scalars of norm one, parametrized by the measurement angles of the input qubits, B_i are signs matrices, depending on the geometry of underlying open graph state, and $\overrightarrow{\varphi}$ is a vector parametrized by the measurement angles of measured auxiliary qubits. The primary structure offers the following simple observations concerning the matrix M:

- The first column is determined only by the geometry of the open graph state and the measurement angles of the auxiliary qubits.
- All the entries of each column are sums of complex numbers of a fixed set, possibly differing in signs.
- The measurement angles of input qubits parametrize the global phase factors of columns of matrix M, which otherwise depend only on the geometry of the open graph state and the measurement angles of the auxiliary qubits.

Moreover we can use this characterisation to easily prove the following simple lemma about uniform determinism. Recall that a pattern is called uniform deterministic if it is deterministic for all possible angles of measurements.

Lemma 19. *A pattern is uniformly deterministic if and only if it is deterministic for all possible choices of auxiliary measurement angles.*

We then proceed to meticulously dissect the B_i matrices to reveal their structure given by the following decomposition:

$$B_i = \gamma_i \Delta_i S B N \Omega_i, \tag{5.2}$$

where γ_i is a sign, which depends on the adjacency of the input qubits, Δ_i, S, N and Ω_i are diagonal sign matrices parameterised by the adjacencies of the set of input to the set of output qubits, the adjacency of output qubits, the adjacency of measured auxiliary qubits, and the adjacency of the set of input to the set of measured auxiliary qubits, respectively. B is a full sign matrix, parametrized by the adjacency between the set of output and the set of measured aux-

iliary qubits. The scalars and the matrices are given in terms of explicit functions on graphs, represented purely graph-theoretically, as adjacency matrices, and as lists of edges. These functions have group-theoretical properties, which we feel could further be utilised to help elucidate the following open problems:

- Simulation of given unitaries directly in the one-way model, *i.e.* without reference to the circuit-based model.
- Characterisation of graph states which implement the same map in the positive branch.
- A refined characterisation of determinism.
- Characterisation of the pointless measurement [92] which is a key element in defining new error correcting codes.

5.2 Preliminaries

In this section we briefly review the one-way model, and present the Phase Map Decomposition [90, 91] of one-way patterns. A brief summary of linear algebra, and the notation used throughout this chapter is given in Section 5.9.

The process of computation in the one-way model can be summarised in the following steps:

1. The setting up of n input qubits in an input state $|\psi\rangle$
2. The addition of $m-n$ auxiliary qubits, prepared in the state $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$
3. The pairwise entanglement of some qubits by means of the $\wedge Z$ interaction. This interaction is represented by an open graph state, an ordered triplet (Γ, I, O), where Γ represents the entanglement graph (two qubits are entangled if and only if the corresponding vertices are adjacent), I is the set of input qubits/vertices and O is the set of output qubits/vertices which is a subset of the auxiliary qubits.
4. The measurement of the input qubits and non-output auxiliary qubits (which we call *pure auxiliary qubits*) in the (X, Y) Bloch sphere plain, that is in the basis pairs $\{(|+_{\alpha_i}\rangle, |-_{\alpha_i}\rangle)\}$, parametrized by a set of measurement angles $\{\alpha_1, \ldots, \alpha_{m-n}\}$. Here we use the following shorthand notation: $|\pm_\alpha\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm e^{i\alpha}|1\rangle)$. The set O corresponds to the qubits which will not be measured.

Without loss of generality we assume input and output qubits are not overlapping. This is not a restriction, as additional auxiliary qubits can be added, which will correspond to the overlapping qubits, to which the quantum state of the overlapping qubits can be teleported. It can be easily shown these two scenarios are equivalent. As quantum measurements are generally probabilistic, the pattern implements a general completely positive map [26]. The scenario in which each measurement corresponds to the projection into the state $|+_{\alpha_i}\rangle$ state is called the positive branch, and the positive branch realises a linear transformation of the Hilbert space of the input qubits to the Hilbert space of the output qubits. The corresponding model is called projection-based quantum computing.

We focus on the positive branch only, for this not to be a restriction, it will suffice that the graph Γ, defined by the underlying graph state, fulfils the graph-theoretical condition of having *flow* or *generalised flow* [35, 39], as then by means of local single qubit corrections, conditioned on sequential measurement outcomes, the entire quantum evolution of the system can be driven to be equal to the positive branch.

We will choose the labelling of qubits so that the first n labels correspond to the input qubits, the following $a = m - 2n$ correspond to the measured auxiliary qubits (which we will call pure auxiliaries), and the last n correspond to the output qubits. We have assumed that in a given one-way pattern all the input qubits are measured first (the first round of the computation). One could easily adapt the whole discussion of this chapter to the scenario where there exist no input qubits or some of the pure auxiliary qubits are also measured in the first round by labelling such qubits among the first n labels.

The measuring of a qubit in the $\{|\pm_\alpha\rangle\}$ basis is equivalent to first locally rotating that same qubit by the local Z_α unitary transformation, followed by a measurement in the $\{|\pm\rangle\}$ basis. For reference reasons, here we give the matrix representations of $\wedge Z$ and Z_α in the computational basis:

$$\wedge Z = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \qquad Z_\alpha = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\alpha} \end{bmatrix} \tag{5.3}$$

Hence, since $\wedge Z$ and Z_α are commuting, the projection-based computation process can be re-stated as follows:

1. The setting up of n input qubits in the input state $|\psi\rangle$
2. The addition of $m - n$ auxiliary qubits, set in the state $|+\rangle$
3. The application of local Z rotations to the input and $m-2n$ auxiliary qubits, corresponding to the measurement angles $\{\alpha_1, \ldots, \alpha_{m-n}\}$
4. The pairwise entanglement of some qubits by means of the $\wedge Z$ interaction. This interaction is represented by an open graph state, an ordered triplet (Γ, I, O), where Γ represents the entanglement graph (two qubits are entangled if and only if the corresponding vertices are adjacent), I is the set of input qubits/vertices and O is an n qubit/vertex subset of the auxiliary qubits representing the output qubits
5. The projection of the input qubits and $m - 2n$ auxiliary qubits to the $|+\rangle$ state

The first and second steps above comprise an embedding of a 2^n dimensional Hilbert space to a 2^m dimensional Hilbert space which we will denote P (for *preparation map*), given explicitly as:

$$P : |\psi\rangle \to |\psi\rangle \otimes |+\rangle^{\otimes(m-n)} \tag{5.4}$$

The application of the $m-n$ local rotations implements a map which we denote Φ_1:

$$\Phi_1 = \prod_{i=1}^{m-n} \mathbf{Z}_{-\alpha_i}^{(i)} \tag{5.5}$$

$\mathbf{Z}_{-\alpha_i}^{(i)}$ denotes an m-qubit unitary, which acts trivially on the composite subspaces of all qubits, except for the i^{th} qubit, where it preforms the $Z_{-\alpha_i}$ rotation. Note that this is an operator on a 2^m dimensional Hilbert space. We collect the entangling interactions, $\wedge Z$, into the map Φ_2:

$$\Phi_2 = \prod_{(i,j)\in\mathcal{E}} \wedge\mathbf{Z}_{i,j} \tag{5.6}$$

where the indexing goes across the set of unordered edges $\mathcal{E}$ of the graph state given by the graph Γ:

$$\mathcal{E} = \{\{v_i, v_j\} \,|\{v_i, v_j\} \subseteq V(\Gamma)\} \tag{5.7}$$

The operator $\wedge\mathbf{Z}_{i,j}$ is an m-qubit unitary transformation, which acts trivially on the component subspaces of all qubits, except the composite subspace of qubits i and j, where it preforms the $\wedge Z$ transformation. We call the cumulative action of the latter two maps the *Phase map* and denote it Φ :

$$\Phi = \Phi_2\Phi_1. \tag{5.8}$$

The last step of the computation consists of projecting the first $m-n$ qubits to the state $|+\rangle$, which we denote R (for *restriction map:*)

$$R = \langle+|^{\otimes(m-n)} \otimes I_{2^n} \tag{5.9}$$

where I_{2^n} is the identity map on the 2^n-dimensional Hilbert space.

Now, the entire process of computation in the projection-based model is represented by

$$R\Phi P \tag{5.10}$$

and we call this representation the *Phase map decomposition* of a given unitary operator implemented in the one-way pattern [90, 91]. Note that one can also derive directly a phase may decomposition for any unitary operators without any references to the one-way pattern [90].

5.3 Structural characterisation of the phase map decomposition

Let $\{|i\rangle\}_{i=1}^{2^n}$ denote the standard computational orthonormal basis of a 2^n dimensional complex Hilbert space. Every computational basis in this representation describes a sequence of 0-1 which is the binary representation of the integer value $i-1$. Therefore $i-1$ represented in binary,

encodes the choice of states $|0\rangle$ or $|1\rangle$ in the component single qubit state spaces. For example, the state $|3\rangle$, in a four qubit setting, represents the state $|0\rangle|0\rangle|1\rangle|0\rangle$ as $(3-1) = (0010)_2$.

Next we refine the expression $R\Phi P|i\rangle$ to obtain the structure of the i^{th} column of the matrix which represents $R\Phi P$ in the computational basis.

Theorem 7. *Let $R\Phi P$ be a phase map decomposition corresponding to a positive branch of a one-way pattern over m qubits, with n non-overlapping input and output qubits, $a = m - 2n$ measured auxiliary qubits, with the set of measurement angles $\{-\alpha_1, \ldots, -\alpha_{n+a}\}$. Then, the matrix M representing $R\Phi P$ is characterised with respect to columns by the following equality:*

$$M\mathrm{e}_i = \varepsilon_i B_i \overrightarrow{\varphi} \tag{5.11}$$

where

- e_i *is the i^{th} vector of the canonical basis*
- $\varepsilon_i = \left(\bigotimes_{k=1}^{n} \begin{bmatrix} 1 \\ e^{i\alpha_k} \end{bmatrix}, \mathrm{e}_i \right)$, *with $(\cdot,\cdot)$ denoting the symmetric dot product*
- $\overrightarrow{\varphi} = \bigotimes_{k=n+1}^{n+a} \begin{bmatrix} 1 \\ e^{i\alpha_k} \end{bmatrix}$
- B_i *is a matrix of signs of dimension $2^n \times 2^a$, which depends on the underlying graph state, and we call them the* sign pattern matrices

Proof. The proof is based on simple linear algebra manipulations so we put the details in Section 5.9. The main properties used are the diagonal form of both Z_α and $\wedge Z$ in the computational basis. The complex phases arising form the Z_α local rotations are collected in the $\overrightarrow{\varphi}$ vector and in the scalars ε_i, and the diagonal of the Φ_2 entangling operation gets spread across the sign pattern matrices B_i. The proof itself presents this structure of the B_i matrices (see Section 5.9)

$$B_i = \sum_{j=1}^{2^n} \sum_{l=1}^{2^{(m-2n)}} b_{\left[(i-1)2^{(m-n)}+(l-1)2^n+j\right]} |j\rangle\langle l| \tag{5.12}$$

These properties will be used in the following section. □

For simplicity, in the expression for $\overrightarrow{\varphi}$ as a numerical vector, we omit a normalising factor of $2^{-(\frac{m-2n}{2})}$, along with the scaling factor ($2^{-(m-n)}$) of the B_i matrices as it has no bearing on the structure we wish to present.

A few direct consequences of Theorem 7 are easily checked:

- The first column of M is parametrized by the measurement angles of pure auxiliary qubits only, as $\varepsilon_1 = 1$.

- For every column i, the entries per row, are of the form

$$\varepsilon_i(r, \overrightarrow{\varphi}) \tag{5.13}$$

  for some vector of signs r. So, entries of a column are sums of the same set of elements, varying possibly in sings only.

- From 5.13 it is clear that every entry of every column is a sum of elements of the set of entries of the vector $\overrightarrow{\varphi}$ varying in signs, multiplied by the column's corresponding global phase factor ε_i.

As mentioned before we can also prove the following simple lemma about the uniform determinism.

Lemma 20. *A pattern is uniformly deterministic if and only if it is deterministic for all possible choices of auxiliary measurement angles.*

Proof. Due to Theorem 7, the measurement angles of the input qubits appear only as global phase factors of the columns of M, and these global factors ε_i are of norm one. Hence the choice of measurement angles of input qubits do not influence the norm of the columns. Also regardless of the measurement angles of the input qubits (as the product of two complex numbers of norm 1 is always norm 1), the matrix M is orthogonal since its columns are orthonormal. Therefore, uniform determinism can only depend on the measurement angles of the measured auxiliary qubits.□

The statement of Theorem 7 indicates a direct method for addressing problems of equalities of patterns, and of simulating a given unitary evolution of a quantum system in the one-way model. For the first problem, we have to evaluate and check the equalities of two expressions of the form of the right-hand side of 5.11. However, that entails knowing how to construct the sign pattern matrices from given graph states. This demands further analysis of the sign pattern matrices, which will be the topic of the next two sections.

5.4 Graph-theoretical characterisation of sign pattern matrices

In the proof of the Theorem 7, the matrices B_i were defined as representations of the expression

$$\sum_{j=1}^{2^n} \sum_{l=1}^{2^{(m-2n)}} b_{\left[(i-1)2^{(m-n)}+(l-1)2^n+j\right]} |j\rangle\langle l| \tag{5.14}$$

in the computational basis, and in that representation b_l corresponds to the the l^{th} diagonal entry of the matrix representation of Φ_2 in that same basis. We now link the graph theoretical aspects of the graph state defining the pattern and the above expression.

Recall that the $\wedge Z$ interaction is diagonal in the computational basis, hence the map Φ_2 is diagonal in that basis as well. We introduce the *sign parity* function SP defined as $SP(k) = (-1)^k$

as it visually simplifies the expressions. It was shown in [90] that b_l, the l^{th} diagonal element of Φ_2 is given by the following expression:

$$b_l = SP\left(\sum_{(i,j)\in\mathcal{E}} x_i x_j\right), \tag{5.15}$$

where x_k was defined as the k^{th} most significant digit (k^{th} digit from the left, including leading zeroes) of $l-1$ represented in binary, and $\mathcal{E}$ is an unordered list of edges of the entanglement graph state, represented by the graph Γ. It is easy to give a graph-theoretical representation for the expression for b_l. It can be shown that the expression $\sum_{(i,j)\in\mathcal{E}} x_i x_j$, where x_i and x_j are functions of l, counts the number of edges of a vertex-induced subgraph of the graph Γ, where the vertex set which induces the subgraph is determined by l. In order to explain how l determines a subset of vertices, we define a function which realises this vertex determination process by an integer:

Definition 21. *If V is a set of vertices, labelled with integers $\{1,\ldots,m\}$, and k is an integer in $\{1,\ldots,2^m\}$ then the selection function Sel is defined as follows: $Sel(V,k)$ is a subset of V such that for the vertex labelled with l (which we present in the subscript) $v_l \in V$, $v_l \in Sel(V,k)$ holds if and only if the l^{th} most significant digit of the m digit binary representation (including leading zeroes) of $k-1$ is 1.*

This function easily extends to any finite totally ordered set O, via an order-preserving bijection between O and $\{1,\ldots,|O|\}$. Also, we will use the expression of the form *a subset of S, selected by (the integer) k* to mean precisely $Sel(S,k)$. Using the introduced terminology, b_l is the sign parity of the number of edges of the vertex-induced subgraph of Γ induced by a subset of the vertices of Γ, *selected by* l.

Now, we can direct our attention to the expression 5.12 and state the following proposition about the graph-theoretical characterisation of the sign pattern matrices B_i .

Proposition 22. *Let Γ be the a graph of a graph state, with vertices labelled by integers $\{1,\ldots,m\}$, such that the first n and last n correspond to input and output vertices (qubits) respectively, and the remaining $a = m-2n$ vertices correspond to the measured auxiliary (pure auxiliary) vertices (qubits) then every entry $(B_i)_{p,q}$, is a sign parity of the number of edges of a vertex-induced subgraph of Γ, and the inducing set of vertices V' depends on the triplet (i,p,q) as follows:*

$$V' = Sel(I,i) \cup Sel(Aux,q) \cup Sel(O,p), \tag{5.16}$$

where O denotes the subset of output vertices, Aux denotes the subset of pure auxiliary vertices, and I denotes the subset of input qubits.

Proof.

Recall that matrices B_i contain the diagonal elements of the matrix Φ_2, which we denote b_l, l indexing the diagonal element. They are explicitly given in the proof of Theorem 7, and they can

be written as

$$B_i = \sum_{p=1}^{2^n} \sum_{q=1}^{2^{(m-2n)}} b_{\left[(i-1)2^{(m-n)}+(q-1)2^n+j\right]} \mathrm{e}_p \left(\mathrm{e}_q\right)^\tau , \tag{5.17}$$

where we have substituted the geometric vectors $|q\rangle$ and $|p\rangle$ with their numerical representatives in the computational basis. In [90] it was shown that b_l, the l^{th} diagonal element of Φ_2 is given by the expression:

$$b_l = SP\left(\sum_{(i,j)\in\mathcal{E}} x_i x_j\right), \tag{5.18}$$

which, as we have noted, has the following graph-theoretical interpretation: b_l is the sign parity of the number of edges of the vertex-induced subgraph of the underlying graph, where the vertex-inducing set is selected by l.

From expression 5.17, the (p,q) entry of the matrix B_i is given by $b_{\left[(i-1)2^{(m-n)}+(q-1)2^n+p\right]}$, i.e. it is the sign parity of the number of edges of the vertex-induced subgraph of the underlying graph, where the inducing subset of vertices is selected by the index $l = \left[(i-1)2^{(m-n)} + (q-1)2^n + p\right]$. The selected set of vertices is easily recognised by observing the binary representation of $(k-1)$:

$$(l-1)_2 = (\overbrace{d_1,\ldots,d_n}^{(i-1)_2},\overbrace{d_{n+1},\ldots,d_{m-n}}^{(q-1)_2},\overbrace{d_{m-n+1},\ldots,d_m}^{(p-1)_2})$$

where d's are binary digits. Recall that in this chapter the input vertices are labelled $\{1,\ldots,n\}$, the pure auxiliary $\{n+1,\ldots,m-n\}$ and the output vertices $\{m-n+1,\ldots,m\}$. Hence, i (the choice of matrix B_i) selects a subset of input vertices. The choice of p (the row of B_i) selects a subset of the output vertices. And q (the choice of column of B_i) selects a subset of pure auxiliary vertices. It follows that $(B_i)_{p,q}$ is the sign parity of the number of edges of the vertex-induced subgraph of the underlying graph, where the inducing set is the union of the subsets of input, output and auxiliary vertices, as chosen by i, p and q respectively, and the proposition is proved.□

Using the terminology of the selection function we can restate this proposition in the following fashion. The (p,q) entry of the sign pattern matrix B_i is the sign parity of the number of edges of a vertex induced subgraph of Γ. This inducing subset is a union of subsets of the input, output and pure auxiliary vertices. The index of B_i (the corresponding column of M) i selects a subset of the input vertices. The row p selects a subset of the output vertices. Finally, the column q selects a subset of the pure auxiliary vertices.

Example We illustrate the application of Theorem 7 and Proposition 22 on a simple but non-trivial underlying graph state: the 2×3 cluster state. This cluster state is illustrated in Figure

5.4. The leftmost pair of qubits (labeled 1 and 2) are the input qubits, and the rightmost (labeled 5 and 6) are the outputs. We denote the local rotation angles (corresponding to the measurement angles) with $\{\alpha_1, \ldots, \alpha_4\}$. The vertex labels are given lower left of the vertices.

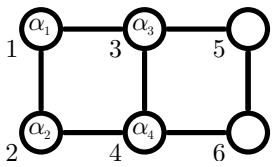

Figure 5.1. The 2×3 *cluster state*

The table given in Figure 5.2 illustrates the application of Proposition 22 on this cluster state, by calculating the corresponding matrix B_k. The rows of the table correspond to the rows of B_4, and the columns to the columns of B_4. The entries of the table are color coded. The red represents the data associated with the choice of the matrix B_4. The blue the data associated with the choice of row of B_4, i.e. p. And the green is associated to the choice of the column of B_4. Each entry of the table contains a cluster state with the vertex-inducing set emphasised by the encircling of the qubits. The final entry of the matrix B_4 are given as the boxed values.

In summary, the matrix B_4 is given as

$$B_4 = \begin{bmatrix} \text{-1} & 1 & 1 & 1 \\ \text{-1} & \text{-1} & 1 & \text{-1} \\ \text{-1} & 1 & \text{-1} & \text{-1} \\ 1 & 1 & 1 & \text{-1} \end{bmatrix}$$

By Theorem 7 fourth column of the matrix M representing the positive branch of a computation preformed on this cluster state can now be given as:

$$M\mathbf{e}_4 = \varepsilon_4 B_4 \overrightarrow{\varphi}.$$

Note that $\varepsilon_4 = e^{i(\alpha_1+\alpha_2)}$, and $\overrightarrow{\varphi} = \begin{bmatrix} 1 \\ e^{i\alpha_3} \end{bmatrix} \otimes \begin{bmatrix} 1 \\ e^{i\alpha_4} \end{bmatrix}$, so the column is characterised by the following expression:

$$M\mathbf{e}_4 = e^{i(\alpha_1+\alpha_2)} \begin{bmatrix} \text{-1} & 1 & 1 & 1 \\ \text{-1} & \text{-1} & 1 & \text{-1} \\ \text{-1} & 1 & \text{-1} & \text{-1} \\ 1 & 1 & 1 & \text{-1} \end{bmatrix} \begin{bmatrix} 1 \\ e^{i\alpha_4} \\ e^{i\alpha_3} \\ e^{i\alpha_3+\alpha_4} \end{bmatrix}.$$

Theorem 7 and Proposition 22 could be potentially used in address the following problems. The equality of patterns and the simulation of a given unitary. In doing so, the essential expression we need to calculate is the expression 5.13. If we are interested in verifying the equality of two patterns, we need to calculate and compare the matrices of their phase map decompositions, given by the Theorem 7. This entails calculating the dot product of the rows of the matrices B_i and the vectors $\overrightarrow{\varphi}$. Similarly, if we are trying to simulate a given unitary, expressions 5.13 which will

B_i, $i=4$ $(1\,1) \to \{1,2\}$	q (column)	1	2	3	4
p (row)	$[q]_2 \to Sel(Aux, q)$ / $[p]_2 \to Sel(O, p)$	$(0\,0) \to \{\}$	$(0\,1) \to \{4\}$	$(1\,0) \to \{3\}$	$(1\,1) \to \{3,4\}$
1	$(0\,0) \to \{\}$	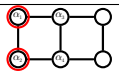 $(1\,1\,0\,0\,0\,0)_2$ $\#E = 1$ $SP = \boxed{-1}$	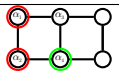 $(1\,1\,0\,1\,0\,0)_2$ $\#E = 2$ $SP = \boxed{1}$	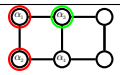 $(1\,1\,1\,0\,0\,0)_2$ $\#E = 2$ $SP = \boxed{1}$	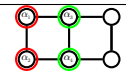 $(1\,1\,1\,1\,0\,0)_2$ $\#E = 4$ $SP = \boxed{1}$
2	$(0\,1) \to \{6\}$	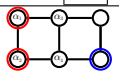 $(1\,1\,0\,0\,0\,1)_2$ $\#E = 1$ $SP = \boxed{-1}$	 $(1\,1\,0\,1\,0\,1)_2$ $\#E = 3$ $SP = \boxed{-1}$	 $(1\,1\,1\,0\,0\,1)_2$ $\#E = 2$ $SP = \boxed{1}$	 $(1\,1\,1\,1\,0\,1)_2$ $\#E = 5$ $SP = \boxed{-1}$
3	$(1\,0) \to \{5\}$	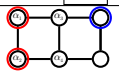 $(1\,1\,0\,0\,1\,0)_2$ $\#E = 1$ $SP = \boxed{-1}$	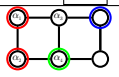 $(1\,1\,0\,1\,1\,0)_2$ $\#E = 2$ $SP = \boxed{1}$	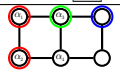 $(1\,1\,1\,0\,1\,0)_2$ $\#E = 3$ $SP = \boxed{-1}$	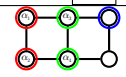 $(1\,1\,1\,1\,1\,0)_2$ $\#E = 5$ $SP = \boxed{-1}$
4	$(1\,1) \to \{5,6\}$	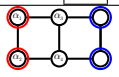 $(1\,1\,0\,0\,1\,1)_2$ $\#E = 2$ $SP = \boxed{1}$	 $(1\,1\,0\,1\,1\,1)_2$ $\#E = 4$ $SP = \boxed{1}$	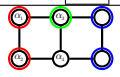 $(1\,1\,1\,0\,1\,1)_2$ $\#E = 4$ $SP = \boxed{1}$	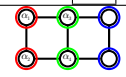 $(1\,1\,1\,1\,1\,1)_2$ $\#E = 7$ $SP = \boxed{-1}$

Figure 5.2. Calculation of matrix B_4 for the 2×3 cluster state

contain variables, as we go across all entries of all columns of the matrix M, will form the left-hand sides of a system of equations we will have to solve (the right-hand side being the entries of the given unitary).

The dot product of rows of the sign pattern matrices and the vector $\overrightarrow{\varphi}$ is in general hard to evaluate, as both vectors have an exponential lengths in the number of pure auxiliary qubits. However, the vector $\overrightarrow{\varphi}$ is represented as a Kronecker product of vectors of length 2, as it corresponds to a state space vector which can be represented as a tensor product of the minimal, two-dimensional component spaces. Such a representation contains the same number of 2-dimensional vectors, as there are measured auxiliary qubits, and so is efficient. The ability to represent the rows of the sign pattern matrices in such a compact form might assist in deriving techniques for solving and evaluating such expressions efficiently.

Hence, in the following section we focus our attention to the structure of rows of the sign pattern matrices and present the decomposition theorem for the sign pattern matrices.

5.5 Decomposition of the sign pattern matrices

If we turn our attention to any row p of any matrix B_i, from Proposition 22 we can see that by selecting the index i (equivalently, a column of M) and a row p we have fixed a subset of the input qubits and a subset of output qubits, respectively. The p^{th} row of B_i is then generated by the sign parities of the numbers of edges of the vertex-induced subgraphs of Γ, where the inducing set is a union of the selected fixed sets of input and output vertices, and the entry of that row (the column of B_i) then selects the additional subset of the pure auxiliary vertices.

Therefore, for fixed p and i, the corresponding row of B_i, which we denote by r, is given entry-

wise by the following expression:

$$(r)_k = SP\left(\#E\left(\Gamma_{Sel(I,i)\cup Sel(O,p)\cup Sel(Aux,k)}\right)\right) \tag{5.19}$$

where $\#E(\Gamma)$ denotes the number of edges of the graph Γ, and for a given graph Γ over the set of vertices V, and $V' \subseteq V$, $\Gamma_{V'}$ denotes a vertex-induced subgraph of the graph Γ induced by the set V'. In expression 5.19 only the subsets of pure auxiliary vertices change as we go across the entries of r.

The subgraph inducing vertex subset is expressed as a union of three subsets, two constant, and one variable. Let us denote $A = Sel(O,p)$, $B = Sel(I,i)$ and $X = Sel(Aux,k)$. As we will be dealing with only one graph at a time, we will drop the graph designation and use the shorthand $\#E(A)$ instead of $\#E(\Gamma_A)$. Also, with $\#E(Y \leftrightarrow Z)$ we denote the number of edges joining vertices in Y with vertices in Z in the graph we are observing. It is then easy to see that

$$\begin{aligned}\#E(A\cup X\cup B) = \#E(A) + \#E(X) + \#E(B) + \\ \#E(B\leftrightarrow A) + \#E(A\leftrightarrow X) + \#E(B\leftrightarrow X)\end{aligned} \tag{5.20}$$

Equality 5.20 and the fact that the sign parity function is a homomorphism from additive monoid of integers to the multiplicative monoid of integers ($SP(i+j) = SP(i)SP(j)$) will give a basis for the decomposition of the sign parity matrices. Therefore, we can express the vector r with respect to the entries as follows:

$$\begin{aligned}(r)_k = SP\left(\#E(B)\right) SP\left(\#E(B\leftrightarrow A)\right) SP\left(\#E(A)\right) SP\left(\#E(A\leftrightarrow X)\right)\times \\ SP\left(\#E(X)\,SP\left(\#E(B\leftrightarrow X)\right)\right.\end{aligned} \tag{5.21}$$

Note the dependencies of the factors of the right-hand side of 5.21 with respect to the explicit parameter k of r, parameter p which is the row selection of B_i and parameter i itself which is the choice of the column of M B_i.

1. $SP\left(\#E(B)\right)$ depends on i only, as it corresponds to a choice of the subset of input vertices.
2. $SP\left(\#E(B\leftrightarrow A)\right)$ depends on both i and p, but is independent of k.
3. $SP\left(\#E(A)\right)$ depends on p only as it corresponds to a choice of output vertices.
4. $SP\left(\#E(A\leftrightarrow X)\right)$ depends on p and k.
5. $SP\left(\#E(X)\right)$ depends on k only.
6. $SP\left(\#E(B\leftrightarrow X)\right)$ depends on i and k.

We have represented the fixed row of a sign pattern matrix r by its entries. We will now represent r by using vector functions, defined on graphs, as that will allow for a simple characterisation of B_i matrix entries.

First, we note that, in the list of dependencies of factors which make up an entry of f, the first three are constants in k. The last three factors depend on k, and we shall represent them as components of values (which are vectors) of two different vector functions on graphs attain.

We define a function on simple graphs whose set of vertices are equipped with a strict order.

Definition 23. *Let Γ be a simple graph, where the set of vertices is equipped with a strict order. We define $\mathcal{P}(\Gamma)$ to be a vector of signs of length $2^{|V|}$ given by the following components*

$$\left(\mathcal{P}(\Gamma)\right)_k = SP\left(\#E(Sel(V,k))\right) \tag{5.22}$$

for all $k = 1, \ldots, 2^{|V|}$.

Since we will often be expressing the $\mathcal{P}$ function of some vertex-induced subgraph of a graph, it is convenient to adopt a shorthand notation. If the graph Γ, which we talk about is clear, and S is a subset of its set of vertices, then $\mathcal{P}(S)$ will be shorthand for $\mathcal{P}(\Gamma_S)$. Recall that Γ_S denotes the vertex-induced subgraph of the graph Γ, induced by the set of vertices S.

The other useful function is defined on bipartite graphs.

Definition 24. *Let Γ be a bipartite graph with partitions V and W, where the set of vertices is equipped with a strict order. We define $\mathcal{B}_\Gamma(V,W)$ to be a vector of signs of length $2^{|W|}$, given by the following components*

$$\left(\mathcal{B}_\Gamma(V,W)\right)_k = SP\left(\#E(V \cup Sel(W,k))\right) \tag{5.23}$$

for $k = 1, \ldots, 2^{|W|}$.

Again, if the graph Γ is clear from context, we will omit the subscript Γ, and simply write $\mathcal{B}(V,W)$ instead of $\mathcal{B}_\Gamma(V,W)$. These two functions can be explicitly defined on different representations of graphs, and these representations have potentially useful properties. We give these properties after we have given the theorem about the decomposition of the sign pattern matrices.

The row r can now be expressed (as its transpose, that is as a column) using $\mathcal{B}$ and $\mathcal{P}$ functions. As the goal is to represent a general column r (that is, any of the rows of any matrix B_i), we introduce these parameters for row r - its row index p, and its sign pattern matrix denoted by i. Therefore, $r_{p,i}$ is now expressed as:

$$r_{p,i} = \gamma_i \, c^1_{p,i} \, c^2_p \cdot \left(\mathcal{B}(Sel(O,p), Aux) \odot \mathcal{P}(Aux) \odot \mathcal{B}(Sel(I,i), Aux)\right) \tag{5.24}$$

where I, Aux and O denote the sets of input, pure auxiliary and output vertices and $\odot$ denotes the pointwise product. The order of the components corresponds to the order of factors in the entry-wise representation of r in 5.21.

In 5.24 γ_i is a scalar, corresponding to the factor $SP\left(\#E(B)\right)$ in 5.21. So we can represent it

using the $\mathcal{P}$ function as

$$\gamma_i = (\mathcal{P}(I))_i . \tag{5.25}$$

Also, $c^1_{p,i}$ is a constant scalar for every entry of a fixed row (hence depends on the row, and the choice of B_i), and corresponds to the expression $SP\left(\#E(B \leftrightarrow A)\right)$ and it can be represented using the $\mathcal{B}$ function

$$c^1_{p,i} = (\mathcal{B}(Sel(I,i),O))_p \tag{5.26}$$

Finally, c^2_p is a constant scalar for a fixed row, and does not depend on the choice of B_i, and corresponds to the term $SP\left(\#E(A)\right)$. It can be represented using the function $\mathcal{P}$

$$c^2_p = (\mathcal{P}(O))_p . \tag{5.27}$$

The three row-wise constant factors have been defined as components of vectors which depend on i or are constant. Then, by collecting the components across rows, and indexes i we can easily note the following deconstruction of the sign pattern matrices.

Theorem 8. *Let $V = I \cup Aux \cup O$ be the set of vertices of the graph Γ, tri-partitioned into input, auxiliary and output vertices. Let*

- $\gamma_i = (\mathcal{P}(\Gamma_I))_i$
- $\Delta_i = diag(\mathcal{B}(Sel(I,i),O))$
- $S = diag\left(\mathcal{P}(O)\right)$
- $B = \left[\mathcal{B}(Sel(O,1),Aux),\ldots,\mathcal{B}(Sel(O,2^{|O|}),Aux)\right]^\tau$
- $N = diag(\mathcal{P}(Aux))$ *and*
- $\Omega_i = diag(\mathcal{B}(Sel(I,i),Aux))$

then

$$B_i = \gamma_i \Delta_i S B N \Omega_i. \tag{5.28}$$

Proof. The origin of γ_i is straightforward and the Δ_i and S matrices are a direct consequence of the $\mathcal{P}$ and $\mathcal{B}$ function representations of the scalars $c^1_{p,i}$ and c^2_p given above.

The B matrix contains the first factor in the brackets in the expression 5.24 in each row, which is constant in i, but variable in row p.

Matrix N is the second factor in brackets in expression 5.24 spread across the diagonal of a diagonal matrix. That factor was constant in p and i and by presenting it as a diagonal matrix which multiplies B from the right, we achieve the pointwise multiplication of each row of B with that factor. Analogous reasoning is used for the matrix Ω_i with the difference that it is variable in i. □

Collecting the results of theorems 7 and 8 we get the following corollary:

Corollary 1. *Using the notation of theorems 7 and 8 the matrix M can be represented with respect to columns as:*

$$M\mathbf{e}_i = \varepsilon_i\gamma_i\Delta_i SBN\Omega_i\overrightarrow{\varphi}.$$

While Theorem 8 completely decomposes the sign pattern matrices, for an actual calculation of the presented decomposition it is useful to have the explicit forms of the functions $\mathcal{B}$ and $\mathcal{P}$. In the following section we present different representations of these functions, and present some of their properties which could be helpful in the application of the decomposition of Corollary 1.

5.6 Explicit representations of the $\mathcal{P}$ and $\mathcal{B}$ functions

The function $\mathcal{B}$ has an elegant representation in terms of the adjacency matrix of the bipartite graph. If $\mathcal{A}$ is the adjacency matrix of a bipartite graph, with partitions V and W, and all the labels of V precede the labels of W, then it is of the block form:

$$\mathcal{A} = \begin{bmatrix} 0 & C \\ C^\tau & 0 \end{bmatrix} \tag{5.29}$$

We now define the function $\hat{\mathcal{B}} : \{0,1\}^n \to \{-1,1\}^{2^n}$, such that

$$\hat{\mathcal{B}}(b_1,\ldots,b_n) = \bigotimes_{i=1}^{n} \begin{bmatrix} 1 \\ SP(b_i) \end{bmatrix}. \tag{5.30}$$

It can be shown that if v is the modulo 2 sum of the columns of C, then

$$\mathcal{B}(V,W) = \hat{\mathcal{B}}(v). \tag{5.31}$$

That is, the $\mathcal{B}$ function can be calculated directly from the adjacency matrix of the bipartite graph in question, by using the $\hat{\mathcal{B}}$ function. Moreover, the $\hat{\mathcal{B}}$ function is a monomorphism from the group $(\{0,1\}^n, \oplus)$ to the group $(\{-1,1\}^{2^n}, \odot)$, where $\oplus$ and $\odot$ represent modulo 2 addition and pointwise multiplication, respectively.

The $\hat{\mathcal{B}}$ representation of $\mathcal{B}$ and the monomorphism property are important as described below. Given the adjacency matrix of the underlying graph we can efficiently compute a polynomial number of entries of the matrix-vector multiplication $B\overrightarrow{\varphi}$ (Corollary 1), even though the mere length of a row of B is exponential in the number of auxiliary qubits. It will suffice to use the representation $\hat{\mathcal{B}}$ given on the right-hand side of 5.30, and $\overrightarrow{\varphi}$ represented in the Kronecker

product form, and use the following property of the scalar product on tensor spaces:

$$(\bigotimes_{i=1}^{n} X_i, \bigotimes_{i=1}^{n} Y_i) = \prod_{i=1}^{n} (X_i, Y_i), \tag{5.32}$$

when X_i and Y_i are of equal dimensions.

The monomorphism property also helps in the scenario where we want to calculate $B\Omega_i\overrightarrow{\varphi}$. Since both the rows of B and the diagonal of Ω_i are represented by the $\mathcal{B}$ functions, and hence by the $\hat{\mathcal{B}}$ functions, due to the monomorphism property, the pointwise product of a row in B and the diagonal of Ω_i is again representable in the $\hat{\mathcal{B}}$ form, which easily reads out of the adjacency matrix, so this becomes efficiently solvable as well.

However, there remains the problem of the matrix N, as what we really wish to calculate is $BN\Omega_i\overrightarrow{\varphi}$, which is represented by the $\mathcal{P}$ function. The $\mathcal{P}$ function results the sign parities of the number of edges of all subgraphs of a given graph as a binary vector.

One way to explicitly represent it is by taking the positive part of the directed adjacency matrix of the given graph Γ. That is, we direct the graph in an arbitrary fashion, and in its directed adjacency matrix (which carries 1 and -1 depending on the direction of the directed edges, of the now directed graph) replace all -1's with zeroes. If $\mathcal{A}$ is that matrix then it can easily be seen that

$$(\mathcal{P}(\Gamma))_i = SP\left((\mathcal{A}\left[i\right]_2, \left[i\right]_2)\right), \tag{5.33}$$

where $\left[i\right]_2$ is the binary representation of $i-1$ given as a vector.

If n is the number of vertices, this representation takes n^2 binary digits on input, as they make up the $\mathcal{A}$ matrix.

An alternative representation uses $\binom{n}{2}$ binary digits in the form of an *edge binary list*, which we now define. Let E be an ordered set of pairs of vertices of Γ such that the label of the first vertex in a pair is strictly smaller than the label of the second, all in all $\binom{n}{2}$ of them, and let E be ordered lexicographically according to edges;

$$E = \{(v_i, v_j) | v_k \in V \ \& \ i < j\}. \tag{5.34}$$

Then, for a given graph Γ, $V = V(\Gamma)$ with $\mathcal{E}$ we denote the binary vector of length $\binom{n}{2}$ such that the i^{th} entry of $\mathcal{E}$ is 1 if the i^{th} pair of vertices of E is adjacent in Γ and 0 otherwise. We call this vector the *edge binary list*. It is easy to see that the edge binary list uniquely characterises a simple graph. If Γ is a graph, and $\mathcal{E} = \left(b_1, \ldots, b_{\binom{n}{2}}\right)$ its *edge binary list* then $\mathcal{P}(\Gamma)$ can be explicitly given as

$$\left(\mathcal{P}(\Gamma)\right)_i = \left(\mathcal{P}\left(b_1, \ldots, b_{\binom{n}{2}}\right)\right)_i = \prod_{k=1}^{\binom{n}{2}} (-1)^{b_k X_1(i,k) X_2(i,k)}, \tag{5.35}$$

where

$$X_1(i,k) = \left\lfloor \frac{i}{2^{f(k)}} \right\rfloor \mod 2, \tag{5.36}$$

and

$$X_2(i,k) = \left\lfloor \frac{i}{2^{\left(k-\binom{f(k)}{2}\right)-1}} \right\rfloor \mod 2, \tag{5.37}$$

with

$$f(k) = \left\lfloor \frac{\sqrt{8k-7}+1}{2} \right\rfloor. \tag{5.38}$$

The unappealing functions X_1 and X_2 can be explained more simply. Let (v_p, v_q) be the k^{th} entry of the set E. Then $X_1(i,k)$ is the q^{th} binary digit of binary represented $i-1$, counting from the least significant digit. With the same notation $X_2(i,k)$ is the p^{th} binary digit of binary represented $i-1$, counting from the least significant digit. This representation, even though seems to be the least elegant has one significant properties. For $\mathcal{P}$ defined on edge binary lists,

$$\mathcal{P} : \{0,1\}^{\binom{n}{2}} \to \{-1,1\}^{2^n} \tag{5.39}$$

is a monomorphism from the group $\left(\{0,1\}^{\binom{n}{2}}, \oplus\right)$ to the group $\left(\{-1,1\}^{2^n}, \odot\right)$ where $\oplus$ denotes pointwise modulo 2 addition, and $\odot$ pointwise multiplication.

How to use this, or any other representation of the $\mathcal{P}$ function to help efficiently evaluate or express $N\overrightarrow{\varphi}$, or $BN\Omega_i\overrightarrow{\varphi}$ in conjunction with the $\hat{\mathcal{B}}$ representation of $\mathcal{B}$ remains an open question.

5.7 The Entanglement Role

In the previous section we have briefly addressed the efficiency of calculating a single entry of the linear map M representing the positive branch of the computation realised by a given one-way computer. Here we explain how the underlying structure of the entanglement effects the efficiency of such computations which is closely linked to the problems of efficient classical simulations of quantum computations and quantum systems in general. To show this, we analyse the computation of a single entry of M thoroughly. Let $m_{i,j}$ be the $(i,j)^{th}$ entry of the matrix M, then from Theorem 1 we have

$$m_{i,j} = \mathbf{e}_i^\tau M \mathbf{e}_j = \mathbf{e}_i^\tau \varepsilon_j \gamma_j \Delta_j SBN\Omega_j \overrightarrow{\varphi}$$

Note that ε_j and γ_j are scalars so the expression above rewrites as

$$m_{i,j} = \varepsilon_j \gamma_j \mathrm{e}_i^\tau \Delta_j S B N \Omega_j \overrightarrow{\varphi}$$

which can be further simplified since Δ_j and S are diagonal matrices. Denote the i^{th} diagonal elements of Δ_j and S with δ_{j_i} and s_i, respectively, then

$$m_{i,j} = \varepsilon_j \gamma_j \delta_{j_i} s_i \mathrm{e}_i^\tau B N \Omega_j \overrightarrow{\varphi}$$

The vector-matrix product $\mathrm{e}_i^\tau B$ constitutes the i^{th} row of B, and denoting this row with r_i we get

$$m_{i,j} = \varepsilon_j \gamma_j \delta_{j_i} s_i \left(r_i, N \Omega_j \overrightarrow{\varphi}\right)$$

where $(\cdot, \cdot)$ is the standard inner product.

For a fixed entry (i, j) of M, all the scalars $\varepsilon_j, \gamma_j, \delta_{j_i}, s_i$ are efficiently computable, so for the efficient evaluation of the entry $m_{i,j}$ the computational bottleneck is the evaluation of the inner product

$$\left(r_i, N \Omega_j \overrightarrow{\varphi}\right)$$

The vectors r_i and $N \Omega_j \overrightarrow{\varphi}$ are of length 2^a, where a is the number of pure auxiliary qubits, so direct calculation would be inefficient. However, recall that both r_i and $\overrightarrow{\varphi}$ have a compact form:

$$\begin{array}{rcl} r_i & = & \bigotimes_{k=1}^{a} \begin{bmatrix} 1 \\ SP(b_k) \end{bmatrix} \\ \overrightarrow{\varphi} & = & \bigotimes_{k=n+1}^{a+n} \begin{bmatrix} 1 \\ e^{i\alpha_k} \end{bmatrix} \end{array}$$

where the angles α_k are the measurement angles of the pure auxiliary qubits, and parameters b_k can be read out of the adjacency matrix of the graph, as explained in the previous section. For the purpose of convenience, we briefly summarize how the parameters b_k are calculated form the adjacency matrix of the graph.

The vector of parameters $[b_1, \ldots, b_a]^\tau$ for the i^{th} row of the matrix B which we denoted r_i is obtained as follows. Let $\mathcal{A}$ be the adjacency matrix of the underlying graph of the computation, respecting the labeling of the vertices (i.e. input vertices precede pure auxiliary vertices, which in turn precede the output vertices). Now, let $\mathcal{A}'$ be a submatrix of $\mathcal{A}$ obtained by dropping all the rows of $\mathcal{A}$ except for the rows corresponding to the pure auxiliary vertices, and dropping all the columns except the columns which correspond to the vertices which are in the set $Sel(O, i)$, that is the subset of output vertices, selected by i. The vector $[b_1, \ldots, b_a]^\tau$ is the modulo 2 sum of the columns in the submatrix $\mathcal{A}'$.

Furthermore, the matrix Ω_j is a diagonal matrix, and by applying the reasoning of the previous section, it is easily seen that its diagonal can also be written in a compact tensorial form:

$$\Omega_j = diag\left(\bigotimes_{k=1}^{a} \begin{bmatrix} 1 \\ SP(o_k)) \end{bmatrix}\right)$$

where the parameters o_k are also read out of the adjacency matrix of the graph. Concretely, the vector of parameters $[o_1, \ldots, o_a]^\tau$ is the modulo 2 sum of the columns of the matrix $\mathcal{A}''$. $\mathcal{A}''$ is

again a submatrix of the matrix $\mathcal{A}$ obtained by by dropping all the rows of $\mathcal{A}$ except for the rows corresponding to the pure auxiliary vertices, and dropping all the columns except the columns which correspond to the vertices which are in the set $Sel(I, j)$, that is the subset of input vertices, selected by j.

Continuing the analysis, the matrix Ω_j can be rewritten as:

$$\Omega_j = \bigotimes_{k=1}^{a} \begin{bmatrix} 1 & 0 \\ 0 & SP(o_k) \end{bmatrix}$$

Therefore, the vector $\overrightarrow{\varphi}'_j = \Omega_j \overrightarrow{\varphi}$ also has a compact form

$$\overrightarrow{\varphi}'_j = \bigotimes_{k=1}^{a} \begin{bmatrix} 1 \\ SP(o_k)e^{i\alpha_{(k+n)}} \end{bmatrix}$$

by the mixed-product property of the Kroenecker product. Hence, the inner product $(r_i, N\Omega_j \overrightarrow{\varphi})$ can be rewritten as

$$\left(r_i, N\overrightarrow{\varphi}'_j\right)$$

where both r_i and $\overrightarrow{\varphi}'_j$ have a compact representation.

The above analysis highlights the role of matrix N, that depends on the graph of the underlying entangled state, in the efficiency of computing m_{ij}. Now we can illustrate a couple of classes of entanglement graphs for which the calculation of a single entry of the matrix M is a feasible task.

First is the case of graphs for which the induced subgraph of the pure auxiliary qubits is empty. If the pure auxiliary qubits are unconnected (unentangled), the matrix N, of the decomposition of Theorem 8, is the identity matrix. This is easy to see as the diagonal entries of N, in this unconnected scenario, are the sign parities of the number of edges of unconnected subgraphs, hence are all zero. The inner product, which is the bottleneck of the computation of a single entry of M, for these graphs is simply

$$\left(r_i, \overrightarrow{\varphi}'_j\right)$$

Since, both of the vectors in this inner product have a compact representation, the above inner product can be written as

$$\left(\bigotimes_{k=1}^{a} \begin{bmatrix} 1 \\ SP(b_k)) \end{bmatrix}, \bigotimes_{k=1}^{a} \begin{bmatrix} 1 \\ SP(o_k)e^{i\alpha_{(k+n)}} \end{bmatrix}\right)$$

which is equal to

$$\prod_{k=1}^{a} \left(1 + SP(b_k)SP(o_k)e^{i\alpha_{(k+n)}}\right)$$

by the property given in equation 5.32.

The techniques used in the prior example can be generalised to more complicated families of graphs. Let the pure auxiliary vertices be denoted $v_{n+1}, \ldots, v_{n+a}$. For simplicity, assume a is

even. Let the graph underlying the one-way computation be such that the pairs of vertices

$$\left\{\left(v_{n+1+2k}, v_{n+2+2k}\right)\right\}_{k=0}^{a/2}$$

are connected (that is, the pure auxiliary qubits are pairwise entangled, respecting the order of the labels). Then it can be shown that, for this graph, the corresponding matrix N is of the form

$$N = \bigotimes_{k=1}^{(a/2)} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

Then, again using the property given by equation 5.32 and the mixed-product property of the Kroenecker product, the calculation demanding inner product $\left(r_i, N\overrightarrow{\varphi}'_j\right)$ can be rewritten as the inner product of vectors

$$\bigotimes_{k=1}^{(a/2)} \left(\begin{bmatrix} 1 \\ SP(b_{2k-1}) \end{bmatrix} \otimes \begin{bmatrix} 1 \\ SP(b_{2k}) \end{bmatrix} \right)$$

and

$$\bigotimes_{k=1}^{(a/2)} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \left(\begin{bmatrix} 1 \\ SP(o_{2k-1})e^{i\alpha_{(2k+n-1)}} \end{bmatrix} \otimes \begin{bmatrix} 1 \\ SP(o_{2k})e^{i\alpha_{(2k+n)}} \end{bmatrix} \right)$$

which equals

$$\begin{aligned} \prod_{k=1}^{(a/2)} \quad & \{1 + SP(b_{2k})SP(o_{2k})e^{i\alpha_{(2k+n)}} + SP(b_{2k-1})SP(o_{2k-1})e^{i\alpha_{(2k+n-1)}} \\ & -SP(b_{2k})SP(o_{2k})SP(b_{2k-1})SP(o_{2k-1})e^{i(\alpha_{(2k+n-1)}+\alpha_{(2k+n)})}\} \end{aligned}$$

which can again be calculated efficiently.

This observation can be generalized – as long as the largest connected component of the subgraph of auxiliary qubits is at most logarithmic in size, we can, in a brute force way, compute an entry of the map M efficiently. This overlaps with the known results on classical simulability of quantum systems. For instance, in [93] it was shown that any MBQC over a resource state whose Schmidt rank scales logarithmically with the number of subsystems can be efficiently classically simulated. In turn, the Schnidt rank an entanglement monotone [94].

If the pure auxiliary qubits are arbitrarily connected, no general direct method of efficient calculation of entries of M is known and hence, the structure of entanglement of the pure auxiliary qubits (captured above with matrix N) proves to be the bottleneck of the computation of a single entry of M. We plan to investigate further the structure of N to fully characterise the class of one-way patterns with efficient classical simulation.

5.8 Discussion

We have presented a complete structural characterisation of the positive branch of a one-way pattern in terms of its matrix representation in the computational basis. This structure was shown to be intricate and complex yet admitting a high degree of regularity. While it remains unclear how to directly use this regularity to tackle problems such as direct simulation of unitaries in one-way model or full characterisation of pointless measurements and etc., however the proposed structure clearly emphasises the importance of the entanglement. Here, entanglement plays a crucial role in the mathematical structures which arise from mathematical descriptions of the process of quantum computation; If the pure auxiliary qubits are unconnected (unentangled), the matrix N, of the decomposition of Theorem 8, is the identity matrix. In that case, all the entries of the matrix realised by $R\Phi P$ can be quickly evaluated, once an open graph state and the measurement angles are given. If the pure auxiliary qubits are connected, this becomes an exponential task. We get a similar effect if we try to solve one restriction of the problem of simulating a given unitary. In this restricted problem an open graph state is given with the unitary, and it is promised that for a certain choice of angles, the positive branch will implement that unitary. For this promise problem it can be shown that it is easily and efficiently solvable if the pure auxiliary vertices of the given graph are unconnected, for some families of graphs. Clearly, entanglement is again crucial. It is our belief that additional work on understanding the algebraic properties of the $\mathcal{P}$ function, that is, of the graph states represented as sign patterns, may yield efficient algorithms for some instances of hard open problems in the one-way model. Such solved instances can benefit the understanding of quantum computation in general.

5.9 Technical details

5.9.1 Summary of notation

Here we present a brief summary of the notation used throughout this chapter. The algebra used is presented in the Dirac notation.

Qubit states A *qubit* is represented by a two-dimensional complex Hilbert space, called the qubit's state space. A *qubit state* is a vector of unit length in the qubit's state space. The *state space of an ensemble of qubits* is represented by the tensor product of the component state spaces, and the *state of an ensemble of qubits* is a vector of unit length in the state space of the ensemble. With $|0\rangle$ and $|1\rangle$ we denote unit orthonormal vectors in the state space of a qubit, and they constitute the *standard computational basis* of a qubit. $|\pm_\alpha\rangle$ denotes a vectors parameterised by the real angle α (and the choice of $+$ or $-$) defined with respect to the computational basis vectors as

$$|\pm_\alpha\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle \pm e^{i\alpha}|1\rangle\right). \tag{5.40}$$

When $\alpha = 0$, we simply write $|\pm\rangle$.

Unitary transformations Z_α denotes a family of *phase shift unitary transformations*, parametrized by the real angle α, represented in the computational basis with the following matrix:

$$Z_\alpha = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\alpha} \end{bmatrix}. \tag{5.41}$$

When the Z_α rotation is applied to the i^{th} qubit of an ensemble of m qubits, the transformation of the state space of the ensemble is denoted with $\mathbf{Z}_\alpha^{(i)}$, which can be given explicitly with

$$\mathbf{Z}_\alpha^{(i)} = I^{\otimes(i-1)} \otimes Z_\alpha \otimes I^{\otimes(m-i).} \tag{5.42}$$

Here, I denotes the identity operator on a single qubit state space. $\wedge Z$ denotes a unitary transformation on the state space of two qubits. In the computational basis it is given by the following matrix:

$$\wedge Z = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}. \tag{5.43}$$

Note that this operator cannot be represented as a tensor product of single qubit transformations. Hence, it can be used to create *entangled states,* which are multi-qubit states which cannot be represented as tensor products of single-qubit states.

When the $\wedge Z$ transformation is applied to the component state subspace of the i^{th} and j^{th} qubit of an ensemble of m qubits, the transformation of the entire ensemble is denoted with $\wedge\mathbf{Z}_{i,j}$. The eigenvectors of the $\wedge\mathbf{Z}_{i,j}$ transformation are the vectors of the computational basis of the ensemble, with eigenvalue -1 if both the i^{th} and j^{th} qubit are in the state $|1\rangle$ and eigenvalue 1 otherwise.

Miscellaneous

- e denotes the basis of the natural logarithm.
- e_i denotes the i^{th} vector of the canonical basis, i.e. a vector with entries 0 everywhere, except a 1 at the i^{th} entry.
- $\otimes$ represents the tensor product. $X^{\otimes n}$ denotes the n-th tensor power of X, explicitly

$$X^{\otimes n} = \overbrace{X \otimes \cdots \otimes X}^{\text{n times}}. \tag{5.44}$$

The tensor product of matrices (and also numerical vectors, as they are isomorphic to single row or column matrices) is called the Kroenecker and defined explicitly as follows:

If A is an m-by-n matrix and B is a p-by-q matrix, then the Kronecker product $A \otimes B$ is the block matrix

$$A \otimes B = \begin{bmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \cdots & a_{mn}B \end{bmatrix}. \tag{5.45}$$

- $SP(n)$, for an integer n denotes the *sign parity* function defined as

$$SP(n) = (-1)^n. \tag{5.46}$$

- X^τ denotes the transpose of the matrix (or vector) X.
- $(\cdot, \cdot)$ denotes the symmetric dot product; If $X = [x_1, \ldots, x_n]^\tau$ and $Y = [y_1, \ldots, y_n]^\tau$ are vectors, then

$$(X, Y) = \sum_{i=1}^{n} x_i y_i. \tag{5.47}$$

- If Γ is a graph, and $A \subseteq V(\Gamma)$ a subset of the vertices of Γ, Γ_A denotes the vertex-induced subgraph of the graph Γ induced by the set of vertices A.
- If Γ is a graph $\#E(\Gamma)$ is the number of its vertices, i.e. $\#E(\Gamma) = |E(\Gamma)|$. If Γ_A is a subgraph of Γ, the graph designation can be dropped and $\#E(A)$ denotes $\#E(\Gamma_A)$. If A and B are disjoint subsets of the vertices of the graph Γ $\#E(A \leftrightarrow B)$ denotes the number of edges connecting the vertices in A to vertices in B in the graph Γ.
- $\oplus$ denotes the modulo 2 addition. If $X = [x_1, \ldots, x_n]^\tau$ and $Y = [y_1, \ldots, y_n]^\tau$ are vectors of integers, then

$$X \oplus Y = [x_1 \oplus y_1, \ldots, x_n \oplus y_n]^\tau. \tag{5.48}$$

- $\odot$ denotes the pointwise vector product; If $X = [x_1, \ldots, x_n]^\tau$ and $Y = [y_1, \ldots, y_n]^\tau$ are vectors, then

$$X \odot Y = [x_1 \odot y_1, \ldots, x_n \odot y_n]^\tau. \tag{5.49}$$

5.9.2 Proof of Theorem 7

Let $R\Phi P$ be the phase map decomposition [90] of the positive branch of a one-way pattern over m qubits, n of which are input, n output, and $a = m - 2n$ are pure auxiliary qubits. Also let $|i\rangle$ be a vector of the standard computational basis. Then we can directly derive the following:

$$\begin{aligned}
R\Phi P|i\rangle &= I_{2^n} R\Phi P|i\rangle \\
&= \textstyle\sum_{j=1}^{2^n} |j\rangle\langle j|R\Phi P|i\rangle \\
&= \textstyle\sum_{j=1}^{2^n} |j\rangle\langle j|R\Phi \left(|i\rangle \otimes |+\rangle^{\otimes m-n}\right) \\
&= \textstyle\sum_{j=1}^{2^n} |j\rangle\langle j|R\Phi_2\Phi_1 \left(|i\rangle \otimes |+\rangle^{\otimes m-n}\right) \\
&= \left(\left(\otimes_{k=1}^{n}\langle +_{\alpha_k}|\right)|i\rangle\right) \textstyle\sum_{j=1}^{2^n} |j\rangle\langle j|R\Phi_2 \left(|i\rangle \otimes \left(\otimes_{k=n+1}^{m-n}|+_{\alpha_k}\rangle\right) \otimes |+\rangle^{\otimes n}\right)
\end{aligned}$$

For clarity reasons we temporarily omit the row-constant scalar $\left(\left(\otimes_{k=1}^{n}\langle +_{\alpha_k}|\right)|i\rangle\right)$

$$\begin{aligned}
&= \textstyle\sum_{j=1}^{2^n} |j\rangle\langle j|R\Phi_2 \left(|i\rangle \otimes \left(\otimes_{k=n+1}^{m-n}|+_{\alpha_k}\rangle\right) \otimes |+\rangle^{\otimes n}\right) \\
&= \textstyle\sum_{j=1}^{2^n} |j\rangle\langle j|\langle +|^{\otimes (m-n)} \otimes I_{2^n}\Phi_2 \left(|i\rangle \otimes \left(\otimes_{k=n+1}^{m-n}|+_{\alpha_k}\rangle\right) \otimes |+\rangle^{\otimes n}\right) \\
&= \textstyle\sum_{j=1}^{2^n} |j\rangle \left(\langle +|^{\otimes (m-n)} \otimes \langle j|\right) \Phi_2 \left(|i\rangle \otimes \left(\otimes_{k=n+1}^{m-n}|+_{\alpha_k}\rangle\right) \otimes |+\rangle^{\otimes n}\right)
\end{aligned}$$

We note that $\Phi_2 = \sum_{l=1}^{2^m} b_l|l\rangle\langle l|$ where b_l is the l^{th} diagonal element of the diagonal matrix Φ_2,

$$
\begin{aligned}
&= \sum_{j=1}^{2^n} |j\rangle \left(\left(\langle +|^{\otimes(m-n)} \otimes \langle j| \right) \sum_{l=1}^{2^m} b_l |l\rangle\langle l| \right) \left(|i\rangle \otimes \left(\otimes_{k=n+1}^{m-n} |+_{\alpha_k}\rangle \right) \otimes |+\rangle^{\otimes n} \right) \\
&= \sum_{j=1}^{2^n} |j\rangle \left(\sum_{l=1}^{2^{(m-n)}} b_{[(l-1)2^n+j]} \langle l|\langle j| \right) \left(|i\rangle \otimes \left(\otimes_{k=n+1}^{m-n} |+_{\alpha_k}\rangle \right) \otimes |+\rangle^{\otimes n} \right) \qquad (5.50) \\
&= \sum_{j=1}^{2^n} |j\rangle \left(\sum_{l=1}^{2^{(m-n)}} b_{[(l-1)2^n+j]} \langle l|\langle j| \left(|i\rangle \otimes \left(\otimes_{k=n+1}^{m-n} |+_{\alpha_k}\rangle \right) \otimes |+\rangle^{\otimes n} \right) \right) \\
&= \sum_{j=1}^{2^n} |j\rangle \left(\sum_{l=1}^{2^{(m-n)}} b_{[(l-1)2^n+j]} \langle l| \left(|i\rangle \otimes \left(\otimes_{k=n+1}^{m-n} |+_{\alpha_k}\rangle \right) \right) \right) \qquad (5.51) \\
&= \sum_{j=1}^{2^n} |j\rangle \left(\sum_{l=1}^{2^{(m-2n)}} b_{\left[(i-1)2^{(m-n)}+(l-1)2^n+j\right]} \langle l| \left(\otimes_{k=n+1}^{m-n} |+_{\alpha_k}\rangle \right) \right) \\
&= \sum_{j=1}^{2^n} \sum_{l=1}^{2^{(m-2n)}} b_{\left[(i-1)2^{(m-n)}+(l-1)2^n+j\right]} |j\rangle\langle l| \left(\otimes_{k=n+1}^{m-n} |+_{\alpha_k}\rangle \right)
\end{aligned}
$$

So now we summarise the entire expression (reintroducing the omitted scalar):

$$
R\Phi P|i\rangle = \left(\left(\bigotimes_{k=1}^{n} \langle +_{\alpha_k}| \right) |i\rangle \right) \sum_{j=1}^{2^n} \sum_{l=1}^{2^{(m-2n)}} b_{\left[(i-1)2^{(m-n)}+(l-1)2^n+j\right]} |j\rangle\langle l| \left(\bigotimes_{k=n+1}^{m-n} |+_{\alpha_k}\rangle \right) \qquad (5.52)
$$

For simplicity, in (8) we omitted a global scaling factor of $2^{-(\frac{m-n}{2})}$, brought about by the scalar products $\langle +|^{\otimes(m-n)}\langle j||l\rangle$ where they are non-zero, and in (9) the global scaling factor $2^{-(\frac{n}{2})}$, caused by the product $\langle j|+\rangle^{\otimes n}$. The overall (omitted) scaling factor is $2^{-\frac{m}{2}}$.

The expression $\left(\left(\bigotimes_{k=1}^{n} \langle +_{\alpha_k}| \right) |i\rangle \right)$ is a scalar which depends on the column i, and we denote it by ε_i, also let

$$
B_i = \sum_{j=1}^{2^n} \sum_{l=1}^{2^{(m-2n)}} b_{\left[(i-1)2^{(m-n)}+(l-1)2^n+j\right]} |j\rangle\langle l|
$$

be a $2^n \times 2^{(m-2n)}$ matrix expressed in the computational basis that depends on the choice of column i with entries in $\{-1, 1\}$. Finally, denote the numeric representation in the computational basis of the vector $\left(\bigotimes_{k=n+1}^{m-n} |+_{\alpha_k}\rangle \right)$ with $\overrightarrow{\varphi}$, which is independ of the choice of the column. It corresponds to the quantum state of the auxiliary qubits after the local Z_α rotations, but before the entanglement procedure. The entire expression can then be rewritten in matrix notation as:

$$
M\mathrm{e}_i = \varepsilon_i B_i \overrightarrow{\varphi}, \qquad (5.53)
$$

where e_i is the i^{th} vector of the canonical basis.□

Part II

Quantum Digital Signatures

Chapter 6

Introduction to quantum digital signatures

Quantum digital signatures allow a party to send a classical message to other parties which cannot be forged or tampered with the authenticity of which, if confirmed by one party will be confirmed by all other parties as well.

6.1 Private channels, message authentication and digital signatures

The lives we live have become de-localized. In a way which was unimaginable less than a century ago we communicate instantaneously, with picture and sound all over the globe. We achieve this by using heterogeneous optical, electrical and radio-wave channels, over complex physical infrastructures from which abstract media like the internet and the Global System for Mobile Communications (GSM) emerge. Communication beyond person-to-person direct contact has become indispensable in our private lives, business, economy and politics.

Given this communication revolution it is then not surprising that the word "cryptography" has rapidly progressed from obscurity to being commonplace in the last 60 years. In many aspects, it is cryptography that makes all the long-distance types of information exchange appropriate substitutes for, the very often impractical, direct person-to-person contact in a secure environment.

In 1984 [86] (perhaps even in 1969 [95]) a new curious flavour of secure information exchange and processing was born – that of quantum cryptography. Then for the first time it was realized certain cryptographic tasks can be performed *better* by exploiting the hidden quantum nature of reality. To better understand how cryptography and quantumness mix, we begin by first briefly presenting some of the central cryptographic concepts with the following caveat. The largest advancement in cryptography, which in some form or other existed for millennia, occurred in the first half of the last century when the concepts like "learning something", "being ignorant", "secure" became formalized. This was achieved in no small part through the seminal work of Claude Shannon [55], the founder of modern information theory. The theory of cryptography is still evolving, even at the conceptual level. For instance, it was only 12 years ago that the crucial notions of composable security were introduced [83], and assuredly the way we think about cryptographic concepts will continue to change. Throughout the introductory part of this work, we will be introducing increasingly more formal notions pertaining to security, which

will ultimately allow us to produce and prove mathematical theorems guaranteeing the formal properties we require from our systems and protocols, provided the conditions of the theorems have been met. Initially, however, we wish to present a very informal, and hopefully intuitive introduction to the basic notions. We begin with private communication, achieved by the famous QKD protocols, and work out way to the central topic of this work – digital signatures, a less well known task in the quantum information community.

Private channels Arguably, the most common flavour of cryptography aims to preserve the privacy of communication between two or more parties, when the physical channels we can use are not under the perfect control of those parties. There is only one method known, which provably ensures perfectly private communication under such conditions: *the one-time pad* [96]. In the one-time pad, two communicating parties A and B can establish "unconditionally secure" or "information-theoretically" secure communication over an untrusted channel provided they initially share a secret message, a "key" guaranteed to be known only to the two parties. Using the key party A can encrypt any message of her choosing, forward it to party B, who decrypts it using the shared key. Provided the key is of, loosely speaking, the same size as the message, it can be shown that no adversary or eavesdropper can learn anything about the actual message. It can be shown that the one-time pad is resource–optimal for the task of secure communication, in terms of the size of the required pre–shared keys[1] [96].

At first glance, it may seem that we are solving a problem by assuming we already have the solution: if we have a means of sharing a secret message (the key) then we have means of sharing a secret message. However, this protocol does bring an important advancement. What it actually ensures is that having shared *some* secret message at a previous moment in time t_0, at any point in the future t_1 we are capable of securely sharing *any* message. For instance, the new message may depend on the events which have transpired in the time interval (t_0, t_1). From the perspective of private channels, the one-time pad allows the extension of their, for some reason, limited duration. Thus, the task secure communication, or realizing private channels, reduces to the problem of *sharing a secret key*.

For this task, the quantum properties of nature have been shown to be extremely useful. In 1984, the protocol referred to as quantum key distribution (QKD), enabling secure distribution of classical keys by means of quantum mechanical properties, was invented [86]. The name of this protocol may be slightly misleading: QKD does not allow for secure exchange of secret keys with no assumptions, but it does allow the exponential, in principle even unlimited, growth of pre-existing secret keys [2].

[1] Technically speaking, the one-time pad is resource optimal (for a message of size n) if all messages of size n are equally likely to appear as the plain-text. If this is not the case, that it, if prior knowledge about the plaint-texts exists, then a protocol which first compresses the plain-text to a shorter message of length $m \leq n$ (which are now (almost) uniformly distributed), followed by a one-time pad of the shorter message, and a reversed process at the receiver's end is optimal. For instance, if it is known that the two possible messages player A may send to player B are "buy Apple stock" or "sell Apple stock", it is not necessary to encrypt a 16 letter message. It will suffice to (publicly) agree that $m =$"0" stands for buy the stock, and $m =$"1" stands for sell the stock, and later on simply encrypt the one bit message m.

[2] In particular, what one assumes for QKD is the existence of a classical authenticated (but not private!) chan-

This turns out to be as good as possible – secure key exchange with no assumptions is known not to be possible, but with the help of QKD, a small amount of the resource of shared secret keys can be greatly expanded.

While the one-time pad encryption allows for private communication, this is not everything we could desire. For instance, one may wish to have a guarantee that the message has not been tampered with. Clearly, one cannot prevent an adversary destroying the message which has been sent (he cuts the optical fibre we are using, for instance), but in the one-time pad encoding, the adversary may tamper with a message in more elaborate ways, giving him an advantage in some cases. To illustrate this a bit we will explain how a one-time pad may work, and how it may fail. Let two parties A and B share two secret and distinct keys with a third party C, and let the secret be an integer, randomly chosen in the set $K = \{0, \dots 99\}$. So party A has $key_A \in K$ and party B has $key_B \in K$ both chosen uniformly at random. Party C has both keys. Then party A can send a message $m_A \in K$ to party C securely, by computing $c_A = key_A + m_A \mod N$ and forwarding c_A to C. C decrypts the message by computing $c_A - key_A$. This simple protocol is the one-time pad, and it can be shown that, by inspecting c_A, the party B cannot learn anything about the message m_A as he knows nothing about the key key_A. The same security holds for party B against an overly curious party A.

Imagine now that party C is an auction house, and the other two parties A (say Alice) and B (Bob) are involved in a blind auction event. Alice and Bob are supposed to submit their offers offer_A and offer_B. The offers are, say in dollars, and it would never be reasonable to pay more than 99 for the item they bid for. The rules are simple: the party who bids more gets the item. But, naturally, both Alice and Bob wish to bid the smallest amount and still get the item. For instance, Alice would bid ideally $\mathsf{offer}_B + 1$. To avoid such events, Alice and Bob send their offers *encrypted* with their secret private keys key_A and key_B, respectively. Thus, the message Bob sends to the auction house is $\mathsf{offer}_B + key_B \mod N$, and Alice cannot learn what offer_B as it was encrypted by the key key_B. Can she still cheat?

For illustration purposes, let us assume Alice knows that Bob is really keen on getting the item, so he will bid more than 50 bucks, but she still has no idea how much more. Then, what clever Alice can do is intercept the encrypted message $\mathsf{offer}_B + key_B \mod N$ and replace it with $\mathsf{offer}_B + key_B - 50 \mod N$, by simply subtracting 50 modulo N. When the auction house decrypts this message Alice tampered with, they will see $\mathsf{offer}_B + key_B - 50 - key_B \mod N = \mathsf{offer}_\mathsf{B} - 50 \mod N$. Since Alice knew Bob wished to offer more than 50, this value is surely below 50. So Alice simply offers 51 dollars, winning always with an offer smaller than what she could have used if she played fair. The property the one-time pad has which malicious Alice abused above is called *malleability*. In general, a cryptographic protocol is malleable if it is possible for the adversary to transform an encrypted text into another encrypted text such that the two decrypt to in a some sense *meaningfully* related decrypted texts.

If the cryptographic system Alice and Bob use for blind auctioning was non-malleable, Alice could still interfere with Bob's message. However, she would have no idea how her tampering

nel between two communicating parties. A classical authenticated channel can be realized from an untrusted one using short pre-shared keys as we explain presently.

will affect the decrypted value – from her point of view the decryption of what Bob actually sent and the decryption of her corrupted version of the secret message, are not related to the modification she caused. Thus, Alice wouldn't know to tamper with Bob's offer in order to help her win. But nonetheless, she *can* tamper with Bob's offer and could either increase her odds of winning, or get Bob in trouble by making it look like he offered more than he can afford.

This is still not fully satisfactory for Bob. What we need is a means of *authenticating* messages, a means of proving the message comes form the honest sender, and that it has not been tampered in transit. Non-malleability and this property of "authenticablility" are closely related. For instance, by the result in [97], which first introduced the concept of *non-malleable cryptography*, a provably secure authentication protocol based on public keys can be derived from any public key non-malleable cryptographic system. We will address the distinction between public and private key cryptography later.

A crucial difference is that while malleability assumes privacy, for authenticity this is not required. As we will show, this causes significant disparities in the resources required for the realization of private versus authentic channels.

Message authentication In the example above, Alice used the malleability of an otherwise information-theoretically secure private communication scheme, to modify Bob's message and win in a blind auction event. Such a malevolent action would not be possible if there existed a mechanism, which allowed the recipient to check whether the message he is receiving indeed comes from the true sender. More generally, such a mechanism needs to ensure that the message sent from a true sender was not tampered with by an unauthorized party.

In message authentication, the goal is to prevent message tampering, or forging, but the secrecy of the message is not enforced. For instance, in electronic bank transactions it is of far greater importance that the transaction gets authorized only if it is requested by a true owner of the account, than that no-one can tell what transaction is taking place.

In their seminal work, Wegman and Carter [98] demonstrated that authenticated message sending is possible under the assumption that the communicating parties share a secret key. While this very much resembles the setting for the information-theoretically secure private communication using the one-time pad, in contrast to it, the Wegman-Carter construction shows that the secret key required to authenticate a message may be significantly shorter than the message itself.

We will very briefly, and rather informally, present the Wegman-Carter construction, which is based on *universal hashing functions*. A hash function h is a many-to-one function with a finite domain and range. For the convenience of presentation, we may assume the domain is the set of m bit binary strings $M = \{0,1\}^m$ (we call the elements of M *messages*) and the range is a set of t bit strings $T = \{0,1\}^t$ (the elements of which we call *tags*) where $t < m$. Then, we can define an indexed family of *universal hash functions* $H = \{h_i | i \in I \text{ and } h_i \text{ is a hash function}\}$, which have the following property: For every distinct pair of of messages $x, y \in M$, the probability they attain the same tag with respect to h_i when $i \in I$ is chosen uniformly at random, is $1/2^t$.

Formally

$$\forall x, y \in M \wedge x \neq y,\ P_{i \in_R I}\left[h_i(x) = h_i(y)\right] \leq \frac{1}{2^t}. \tag{6.1}$$

In other words two tags of two distinct messages collide with probability at most $1/2^t$ if the hash function is drawn randomly from H.

Note that this is a property characterising the entire set H, and not individual functions h_i. In their work, Wegman and Carter have shown that not only universal families of hash functions exist for all $0 < t < m$, but also that there exists a randomized algorithm which can, for a family H, efficiently generate a description of a function $h \in H$ chosen uniformly at random. Now, we can build a simple Wegman-Carter authentication scheme.

Two parties, Alice and Bob who wish to communicate in an authenticated way, using their pre-shared short key of length t, generate identical hash functions h_{key}. Since Alice and Bob are the only ones who knew the shared key, for the rest of the world h_{key} looks as chosen uniformly at random from a publicly known family H. For the family of universal hash functions they consider, their domain is of the length $m \gg t$ of the message Alice wishes to send to Bob, and the range contains t bit strings. To send a message $text$ to Bob in an authenticated way, Alice computes the t bit *tag* of the message $text$ by using h_{key}, $tag_A = h_{key}(text)$. She appends the tag to the message and sends to Bob. Bob then receives $(text', tag'_A)$ as a third party might have interfered with the message. He then computes the tag of the received message himself, $tag_B = h_{key}(text')$ and accepts Alice's message as authentic only if $tag_B = tag'_A$.

Clearly, unaltered messages of Alice will always be confirmed by Bob. An adversary may attempt to pretend to be Alice and send a message in her name, or modify hers. But, to pass the authentication, he has to produce a valid pair $(text', h_{key}(text'))$. Since he has no information about key, he can pick a random hash function from the publicly known universal set, but by the defining properties of universal hash functions, the probability he generates the correct tag is bounded by $1/2^t$, thus decaying exponentially quickly.

Note that the message itself may be very large, say gigabytes, and the tag (and the pre-shared key) can be very small and still guarantee sufficient security. By sharing say 512 random bits, the probability their message will be successfully tampered with, or forged, will be on the order of 10^{-154}. The security guarantees of this scheme rely only on the mathematically provable properties of sets of universal hash functions and in no way on assumptions on the powers of the adversary – this scheme is information-theoretically secure.

We have mentioned that the QKD protocol is actually a *key expansion* rather than a key generation protocol. QKD comprises two conceptually differing parts. First is a quantum part in which quantum states are exchanged between the two parties and measurements are performed. The second part is a classical post-processing part in which secure identical keys are generated starting from "raw keys". In this part, Alice and Bob diligently communicate in order to carry out *information reconciliation* – guaranteeing that in the end Alice and Bob share identical keys

– and *privacy amplification* [3], in which any residual correlations between Alice and Bob's shared keys and the system of the environment, which may be the eavesdropper, are expunged.

QKD is proven secure if *and only if* the communication between Alice and Bob in the post-processing phase is authenticated. To ensure this, Alice and Bob must employ some type of an authentication scheme, and as we have seen this requires short pre-shared keys. In a nut-shell, short keys allow for authentication of long messages – message authentication used within QKD allows for secret key exchange – secret keys guarantee information-theoretically secure private communication channels. Note that the causal link above makes sense in practice because *authenticated* channels are **substantially cheaper** than *private* channels in terms of the resource of pre-shared secret keys.

Thus, in principle, by sharing a short initial key with Bob, Alice can send a message to Bob of a large size which cannot be tampered with. However, the story gets a bit more complicated if a third party, Charlie, is introduced into the picture. Consider the following prototypical scenario. Alice wishes to buy a house from Bob. Her money is safely stored in a bank ran by Charlie. Bob is happy to sign the deed to the house over to Alice, provided she hands him over a signed contract, a cheque, which Bob can take to Charlie (the bank) and collect his money. Can we resolve this problem using techniques for authenticated messages, and without having Alice go to the bank with Bob?

Digital signatures Using authentication, Alice can prove to Bob the message came from her, and Bob is guaranteed it has not been corrupted in transit. We could try to resolve our three party problem by having Alice send Bob an authenticated message giving him powers to take out her money from Charlie the Bank, but Bob alone can verify it is legitimate, because the keys are shared between Bob and Alice alone. The key could not have been shared between all three of the parties, because then Bob could send a message to the bank *pretending* to be Alice, without striking a deal with her, and get her money without handing over the house.

We can isolate the properties we require to resolve the three party problem above, which is not solved by authentic channels alone. First, we require that every party can confirm (or refute) the authenticity of Alice's message, and that no party will accept a message which was tampered with or not sent by Alice. This property we will call *non-forgeability*, and primarily protects the sender Alice. This could be achieved with multiple authentic channels. However, we require that parties will be consistent in their decisions on authenticity over the same message. This property is called *non-repudiability* and it implies that if Bob authenticates a message from Alice, so will the bank, and it serves to protect the recipients form a dishonest Alice. Without this property, Alice could for instance cheat Bob out of his house! If Alice can repudiate her messages, Alice could sent Bob a contract, which he checks, confirms it came from Alice, and gives Alice keys to the house. Later, he goes to the bank to take out his money. If Alice can cause the bank not to authenticate her contract, then Bob could lose the house and fail to get his money! Thus, the two properties we care about are: 1) No forging – in a multi party setting, no party can send a message to any

[3] Incidentally, privacy amplification is often based on hashing techniques.

other party pretending to be someone else without getting caught. Also, sent messages cannot be tampered with without getting caught. 2) No repudiation – a message confirmed by one party to be authentic will be confirmed to be authentic by all parties as well.

In the olden days, this problem was solved by hand-written signatures, or perhaps by handwriting comparison. In the modern world, the protocol which is used is *digital signatures*. Early protocols for digital signatures involved cryptographic one-way functions (such as the RSA encryption function), and the signing and verification processes closely resembled encryption and decryption processes in public-key cryptographic schemes. Due to their importance in digital economies, digital signature protocols have been a prosperous field in cryptography. One particular type of an early signature scheme invented by Lamport we shall explore in the sections to come. Most practical and used digital signatures schemes nowadays offer high efficiency, low resource requirements, and security guaranteed under computational assumptions. The schemes we present now, on the other hand guarantee information-theoretic security under assumptions of validity of quantum mechanics.

Before we begin with the introduction to the protocol of quantum digital signatures we wish to briefly explain some important flavours of security. In this work we will often contrast *security under computational assumptions* (*computational security*) versus *unconditional* (*information-theoretic*) security. As the names indicate, in computational security, a protocol is secure provided the computational powers of the adversary are in some way limited. More often than not, this entails that to break a code an adversary would have to compute a function, for which no known polynomial algorithm exists on a classical computer. However, there are also schemes which are believed to be secure against quantum computers as well (but not against *computationally unbounded* adversaries).

Concerning *unconditional* security, we point out that one has to be careful how the word unconditionally is understood. No protocol of secret key exchange can, for instance, be secure against an armed adversary who extracts the key from one of the recipients at gun point. Less dramatically, for some of implementations of "unconditionally" secure QKD, teams of professional "quantum hackers" produced attacks, which demolished the claimed security [99, 100]. Their attacks targeted specific implementations of an abstract scheme, from the types of photon detectors used in the implementation which can be manipulated by strong light sources. Different hacking approaches could, for instance, also include radio interference attacks targeting the electronic equipment used. In this sense, one has to distinguish between a *protocol's security* and *the security of implementation*. Proofs of security are formal statements in a formal language of mathematics, and pertain to *models* of systems realizing cryptographic protocols, and can therefore only be proven secure against the types of attacks which can be formulated *in the model*. Models are unlikely to ever describe all possible elements of reality in practice [4] The problems of security of implementation are a very different class of problems from the security of protocols. Thus, when we say unconditional security, we mean the security of the protocol as modelled, and

[4] In a sense, this is the ultimate goal of physics, and perhaps all empirical science. Even if it is possible, a complete description of reality may be a bit cumbersome to work with on paper when security proofs are produced.

not the security of any real implementation.

Nonetheless, information-theoretic security, compared to computational security is a great improvement, provided it does not come at an absurd price (for instance, if we need to generate black holes to ensure a particular protocol works). The issue with computational security is not only that in the future a clever hacker may find ways to compute hard cryptographic functions efficiently, and then break our codes in run-time. The issue is more problematic than that in the sense that *given enough time*, some of the computationally secure codes could be broken with the computers and algorithms we have *now*, by exhaustive search attacks. That is, such cryptography offers security over time-scales which may be large [5], but are *always* limited. The notion of security which takes such considerations into account is sometimes called everlasting security – a protocol is everlastingly secure if it cannot be broken by an adversary that becomes computationally unlimited *after* the protocol execution [101]. Information-theoretically secure protocols are everlastingly secure, whereas many computationally secure protocols do not have to be.

Private-keys versus Public keys[6] In cryptography there are two central approaches how abstract cryptographic functionalities can be implemented. One technique is based on *private* or *symmetric* keys, whereas the other approach on so-called *public* or *asymmetric* keys.

In private key schemes, legitimate parties share a common identical key. The prototypical example of a private key cryptosystems are the one-time pad, enabling secure private communication over untrusted channels. Another is the Wegman-Carter-type authentication protocol, ensuring message authentication over insecure channels. Both examples were described earlier in this chapter. In private-key schemes the same keys are used for data encryption and decryption and the corruption of the secrecy of the private shared key regularly demolishes security. In cryptographic theory, when working in the setting of private keys, the central assumption is that the keys are initially obtained via an undefined process, which is outside the scope of consideration. A frequent characteristic of all private key schemes is that they are information-theoretically secure. In practical cryptography, the assumption of having pre-shared keys is a proverbial "elephant in the room", and while many models for distributing such keys exist, they often include impractical entities, like trusted mobile couriers and trusted centres. Quantum key distribution has, at least in theory, significantly alleviated the pain of the difficult key distribution, by allowing effectively unlimited key expansion. However, as we have mentioned, even QKD requires a certain amount of pre-shared keys to begin with, so while the problem of initial key sharing is reduced in quantity, it is not fully resolved. Indeed in information-theoretically secure settings it cannot be fully resolved[7].

Public-key (or asymmetric key) schemes emerged in 1970's, and their defining property is that

[5] The estimation of the durability of such codes is difficult as it requires an accurate prediction on the advancements in computational technology and theory in years to come.

[6] This section is based on [102, 103]

[7] However, QKD can also be combined with a computationally-secure authentication scheme, in which case, if the authentication scheme was not broken during the key exchange of the QKD protocol, then the generated keys are secure against unbounded adversaries. This was recently proven in [104], and shows that QKD with a computationally-secure authentication can be everlastingly secure.

the encryption and decryption keys differ. Typically in public-key schemes, Alice (a player in a scheme) will produce a pair (public key, private key), emitting the public keys to all interested parties, and keeping the private key to herself. For illustration purposes, secure communication in public-key systems will be achieved by having Bob encrypt his message using Alice's public key, and send it to her. The property of the scheme will be such that only a person with the valid private key can decrypt such a message under reasonable assumptions. Since public-key schemes are asymmetric for two-way communication, both players will have to produce public-private key pairs. Public-key schemes are most often, if not always, secure under certain assumptions. These assumptions are often based on the allowed computational power of realistic adversaries, and rely on computational hardness assumptions of mathematical problems. On the other hand, a practical advantage of public-key schemes over private-key schemes is in the fact that public keys can safely be reused. Private-key schemes tend not to remain secure if iterated with same private keys. In the theory of public-key cryptography, the initial assumption is that public keys can be "broadcast" to all parties, *i.e.* that the starting point of consideration already includes the valid public-key recipients possessing the key. Without this assumption, public-key cryptography is vulnerable to impersonation, or "key exchange" attacks. In these types of attacks a corrupted party swaps a legitimate public-key with his own, and extracts secret information by impersonation. In practice, secure distribution of public-keys is a big issue, often solved by "bootstrapping" techniques. Trusted third party centres, called *certificate authorities* (CA) embed their public keys in the browsers we use, and using these authenticate legitimate users' public keys. These systems are again vulnerable to impersonation attacks, from interested party to CA, and by tampering with the keys in our browsers, but they become more difficult. In such solutions to initial distribution, a key property used is the reusability of the public keys.

In modern cryptography, the distinction between public and private key schemes is strict. In particular, the defining properties of certain cryptographic protocols are not only the functionality, that is what type of guarantees they ensure for the users, what are they *are meant to do*, but whether the protocol is private or public key.

For instance the modern definition of digital signatures implies that it is a public key protocol, which in way similar to the Wegman-Carter authentication scheme, produces a tag associated to a message which is a signature. Formally:

Definition 25. *[102] A* signature scheme *is a triple* (G, S, V) *of* BPP *algorithms satisfying the following two conditions:*

1. *On input* 1^n*, algorithm* G *(called key-generator) outputs a pair of bit strings.*
2. *For every pair* (s, v) *in the range of* $G(1^n)$*, and for every* $\alpha \in \{0,1\}^*$*, algorithms* S *(signing) and* V *(verification) satisfy:*

$$P(V(v, \alpha, S(s, \alpha)) = 1) = 1, \tag{6.2}$$

where the probability is taken over the internal coin tosses of S *and* V*.*

The integer n *serves as the security parameter of the scheme. Each* (s, v) *in the range of* $G(1^n)$

constitutes a pair of corresponding signing/verification keys.

Then $S(s, \alpha)$ is a signature to the document α produced using the signing key s.

In this type of schemes, a large message gets signed with one key, and gets verified using another. The message length need not depend on the length of the key. In the first type of concrete signatures protocols, each bit of every message will be signed by a key. Following this, we will present the basic idea behind Quantum Digital Signatures (QDS). In QDS, the "public keys" are secret quantum states. Quantum states have substantially different properties than classical information which usually comprises public keys. In particular, unknown states cannot be copied, and indeed this property will play a central role in the security of the scheme. Whether QDS is a legitimate public-key scheme may be a delicate question, but for the purposes of this thesis we will focus on the guaranteed security properties – the functionality of the protocol, under assumptions we will make transparent. Whether and how the assumptions can be justified we will briefly address in Chapter 9.

6.2 From one-way functions to quantum digital signatures

In 1979 Lamport devised a digital signature scheme based on one-way functions [105]. Assume the existence of a (deterministic) function f, which is easy to evaluate on each element x of the domain, yet, given an image $f(x)$, computing an element of the pre-image $y \in f^{-1}(f(x))$[8] is computationally very difficult. Let the domain of f be strings of length L from some finite alphabet, say $A = \{0, \ldots, N-1\}$.

Consider a three party setting, with Alice, Bob and Charlie. To sign a single bit message at some point in the future Alice generates two random strings (private keys) $\overrightarrow{priv_0}$ and $\overrightarrow{priv_1}$ of length L from the alphabet A. Then she computes her public keys $\overrightarrow{pub_0} = f(priv_0)$ and $\overrightarrow{pub_1} = f(priv_1)$, and sends the private keys, in addition to a bit value to both Bob and Charlie. Thus, Charlie and Bob both have $(0, \overrightarrow{pub_0})$ and $(1, \overrightarrow{pub_1})$. The function f is assumed to be public.

To send say a message, say $m = 0$, to Bob at some later point, she reveals her private key corresponding to message bit value m and the message. Thus she sends $(0, \overrightarrow{priv_0})$. Bob computes $f(\overrightarrow{priv_0})$ and accepts the message as authentic if $f(\overrightarrow{priv_0}) = \overrightarrow{pub_0}$.

To transfer the message to Charlie, Bob sends $(0, \overrightarrow{priv_0})$ he received from Alice, and Charlie *verifies* the message if it passes the same test Bob did.

The security of this protocol is easy to illustrate, assuming the hardness of inverting f and authenticity of the communication channels.

No repudiating Successful repudiation means that the same message, sent by Alice, gets confirmed by one recipient and rejected by the other. Assuming Charlie and Bob received identical public keys initially, they will use the same deterministic one-way function to authenticate/verify

[8] The function need not be bijective, so the pre-image may be a set.

the signed bits. Thus the outcomes they obtain will be the same – they either both confirm or reject the message, so repudiation cannot occur. Thus, the only recourse for a malicious Alice is to tamper with the public keys. In public-key cryptography settings it is often assumed the public keys are identical and un-tampered with between all recipient parties, in which case Alice has no way to cause repudiation. Without this assumption, repudiation is prevented by having Bob and Charlie compare the messages, which is easy given they are classical information. In the most general setting Bob and Charlie may require authentic channels to implement comparison, which again can be realized using the digital signatures protocol.

No forging For Bob to forge a message, say $m = 1$ (or to alter the one bit message $m = 0$ Alice sent and Bob intercepted), he would, in our example, have to produce one element of the pre-image of the public key $\overrightarrow{pub}_1$ with respect to the function f. But, by assumption of hardness, this cannot be done. Bob can *guess* a pre-image, in which case he will succeed (in the case of bijective one-way function f) with probability of the order of $1/N^L$. Thus, by choosing L large enough, this probability can efficiently be made arbitrarily small.

This is of course not a formal proof, but it will do for illustrative purposes.

In the first paper on Quantum Digital Signatures, Gottesman and Chuang [2] proposed a quantum version of this protocol, where the one-way function f produced mutually non-orthogonal quantum states from a classical input. Thus, the public keys Alice uses are *quantum states*, and private keys are *full classical descriptions* of these states. The protocol goes as follows:

Protocol 12 QDS (original) [2]

- **Preparation**
  A map $k \to |f_k\rangle$ taking classical messages to quantum states is chosen by all parties. Alice chooses a number of pairs of $L-$bit strings $\{k_0^i, k_1^i\}, 1 \leq i \leq M$ uniformly at random. M is a security parameter. The states $\{|f_{k_0^i}\rangle, |f_{k_1^i}\rangle\}_i$ are distributed amongst a fixed pre-defined number of recipients. Two thresholds c_1 (authentication) and c_2 (verification) are appropriately selected and broadcast to all participants.
- Sending a single bit message b

  1. Alice sends $(b, k_b^1, \ldots, k_b^M)$ over an insecure classical channel.
  2. Each recipient of the signed message checks each of the revealed public keys to verify that $k_b^i \to |f_{k_b^i}\rangle$. Recipient j counts the number of incorrect keys; let this be s_j
  3. Recipient j accepts the message as authentic if $s_j \leq c_1 M$. If $s_j \geq c_2 M$ message is rejected. If $c_2 M \leq s_j \leq c_1 M$, the message is authentic, but may be rejected by other recipients.

Security against forging is information theoretical and certified by quantum mechanics: non-orthogonal states cannot be distinguished with unit probability, thus from observing the received public key a potential forger cannot find out the full description of it. In information-theoretical terms the *accessible information* in the quantum states is upper bounded by the Von Neumann entropy of these states, and this value can be much smaller than the entropy (accessible information) in the classical description of the states.

Security against repudiation is, however, much more difficult in the quantum setting. As we have mentioned, in the classical setting this property is almost trivially satisfied because Bob and Charlie can reliably compare the classical information they get from Alice. Comparing quantum information, general quantum states is a non-trivial task. However, by using a general quantum comparison method, called a SWAP test [2], Bob and Charlie can get certain guarantees that the states they have are close enough, in which case security against repudiation can be ensured, again information-theoretically.

This scheme is easily generalized to many recipients. Alice will send her quantum public keys to all of them, and they will compare them using SWAP tests amongst each other. Given that many copies of the quantum states now exist in the environment, Alice will have to make sure that the states she sends are large and complicated enough so that the difference between the amount of information describing the states is larger than the information accessible to the potential forger, even if he has access to all but one copy of the public keys.

We will present a modification of a particular variant of the QDS protocol using sequences of coherent states as public keys. These public quantum keys we will refer to as *quantum signatures*. There we will further investigate the details of the properties of the QDS protocol. The modification we mention lies in the particular way the quantum state comparison is performed, which is specific to the coherent state setting.

6.3 QDS using coherent light

The idea of using multiple coherent states for "quantum lock and key" schemes, which enable quantum public key types of protocols was proposed in [3]. As we noted, for the purposes of quantum digital signatures, the chief impediment is the difficulty of comparing quantum states– the public keys. In this work they present a general simple scheme for comparing two coherent states, which has a higher success probability than the best general quantum state comparison scheme illustrated and explained in Figure 6.1, using a passive device called a "multiport" and photon detectors.

Concretely, if two coherent states $|\alpha\rangle$ and $|\beta\rangle$ are compared, if the states are not the same, the proposed comparison scheme succeeds in detecting the difference with probability

$$p_{succc} = 1 - e^{-\frac{1}{2}|\alpha-\beta|^2} \tag{6.3}$$

compared to the success probability of the optimal general quantum comparison strategy which succeeds with probability

$$p_{succc} = \frac{1}{2}\left(1 - e^{-|\alpha-\beta|^2}\right). \tag{6.4}$$

Hence, the proposed method achieves better results than the most general disambiguation method always, except in the trivial case, when both the amplitudes of the coherent states are zero. The advantage here is gained from the fact we restrict ourselves on the comparison of coherent states

only, which can be understood as prior knowledge.

The multiport presented in Figure 6.1 compares two states, but the concept is easily generalized by considering an array of mutually linked balanced beamsplitters and an array of detectors. For details we refer the reader to [3].

Aside from being more efficient in comparing coherent states than a universal comparison strategy, multiports can be relatively easily realized experimentally since they only comprise passive linear optical elements. Moreover, unlike most proposed or realized experimental schemes for quantum states comparison, the multiport is a *non-demolition* comparator – if the input states are the same, the multiport signal arms in Figure 6.1 contain unperturbed coherent states.

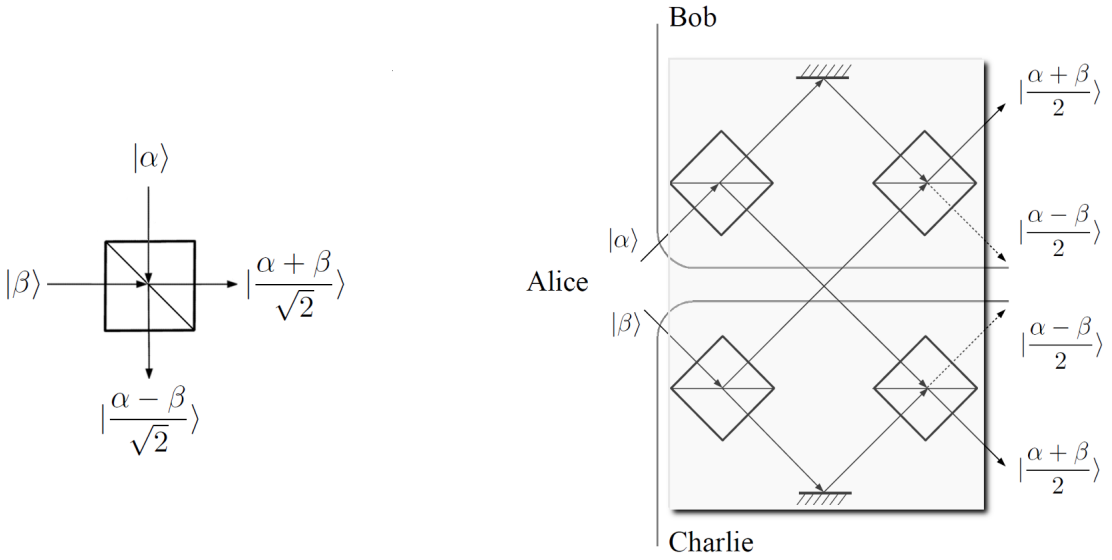

Figure 6.1. Multiport idea. The left-hand side image illustrates how a simple balanced 50:50 beamsplitter can be used to compare coherent states. If the input states are identical, the downward-facing output arm contains the vacuum. A detection of a photon in this "null-port" arm proves the input states were not identical. The right-hand side image illustrates the multiport. The multiport is a passive optical device (in the shaded rectangle) shared by two parties. If Alice introduces unequal states in the multiport, the "null-port arms" containing the $|(\alpha-\beta)/\sqrt{2}\rangle$ states contain a non-vacuum state which can be detected by Bob and/or Charlie. If the states were identical, the non-null-port arms called the signal arms contain the input states. This realizes a type of non-destructive state comparison which does not require the two compared states to be fully in the possession of one party.

6.3.1 A proposal for an experimental realization for a three party quantum signature distribution protocol

Based on the idea for the comparison of coherent states we have just presented the following proposal for a realization of a QDS scheme has been given in [3]. Alice's quantum key will comprise L coherent states, of a fixed mean photon number $|\alpha|^2$:

$$|\psi_{key}\rangle = \bigotimes_{i=0}^{M-1} |e^{I\theta_i}\alpha\rangle \tag{6.5}$$

where $|e^{I\theta_i}\theta\rangle$ may be one of N complex phase shifted coherent states (with respect to a fixed oscillator – reference frame, the phase is defined and α assumed to be real and positive) parametrized by a randomly chosen angle $\theta_i \in \{\frac{k\pi}{N} | k = 0, \ldots, N-1\}$. The parameters $N, L, |\alpha|^2$ are the security parameters of the protocol and the angles $\{\theta_i\}$ will comprise Alice's private key.

Alice will sequentially emit a copy of the component state $|e^{I\theta_i}\alpha\rangle$ to Bob and Charlie who will compare the inbound states using beamsplitters, as it is explained in Figure 6.1.

The states coming out of the signal arms of the multiport can then be used as Alice's quantum signatures, provided the null port arms detected no photons. This holds in the case of ideal devices, however, in the analyses to come imperfections will be taken into account.

This proposal has been realized experimentally, and the results of the experiment are the topic of the following chapters.

Chapter 7

Experimental set-up

We describe the experimental set up designed according the proposal discussed in the previous chapter. Experiments performed relevant for the security analysis are explained.

The proposal for the experiment described in the last section was implemented at Heriot-Watt University, by Patrick J. Clarke and Dr. Robert J. Collins and under supervision of Prof. Gerald S. Buller and Dr. Erika Andersson.

7.1 Basic properties of the experimental system

Before we give details of the experimental set up, we will first give the description of the protocol Alice, Bob and Charlie will run. Then, we will elaborate on how particular elements of the protocol were experimentally realized with additional technical details.

The protocol we present in Protocol 13 distributes one pair of quantum signatures, and performs the authentication and verification for those states. Thus, it enables the signing and verifying "one-half" of a bit. To sign a bit, the distribution would be run twice, and only one of the distributed quantum signatures would be checked against classical information revealed by Alice.

The main conceptual components of the protocol are:

> *Shared:* a medium over which the quantum states can be sent, a common reference frame for the complex phase, and a classical channel which can carry the classical information about the phases – the private keys needed in the signing phase.
>
> *Alice's:* a phase-controlled coherent state generator, and a random number generator.
>
> *Shared by Bob and Charlie:* the multiport.
>
> *Held by both Bob and Charlie:* devices performing verification/authentication based on classical information and the received quantum states.

Additionally, Bob and Charlie would require a quantum memory, but as this technology is not yet mature enough for the purposes of QDS, in our experiment the signing, and authentication/verification is performed in run-time.

Protocol 13 Quantum Signatures Distribution with Authentication and Verification

1. To sign half of a bit, message $m = 0$ in the future, Alice generates a sequence $PrivKey_0 = (\theta_1^0, \ldots \theta_L^0)$ and $PrivKey_1 = (\theta_1^1, \ldots \theta_L^1)$ of L randomly chosen angles from the set of N equally spaced phases, so

$$\theta_k^m \in \{\frac{2r\pi}{N} r = 0, \ldots N-1\}. \tag{7.1}$$

   The pair $(0, PrivKey_0)$ is called a private key pair for message $m = 0$.

2. Alice then generates two copies of a sequence of coherent states $QuantSig_0 = (\rho_1^0, \ldots, \rho_L^0)$ with the coherent phases matching the angles in the sequence $PrivKey_0$, thus $\rho_k^0 = |e^{i\theta_k^0}\rangle\langle e^{i\theta_k^0}|$ where α is a real positive amplitude. A sequence of such states is called a quantum signature. She sends a copy of the quantum signature to each of Bob and Charlie each, informing them that they correspond to message $m = 0$. The individual state ρ_k^m we refer to as the k^{th} quantum signature element state for message m.

3. The sequences of coherent pulses $QuantSig_0$ traverse the multiport, shared by Bob and Charlie, and the output states exiting are stored in quantum memories of the two recipients. The exit null-ports of the multiport are equipped with photon detectors. Photon detection events are tracked by Bob and Charlie.

4. To sign the message $m = 0$ with Bob, Alice announces the message $m = 0$ and the corresponding private key to Bob (thus she sends the pair $(0, PrivKey_0)$ over an untrusted channel). To authenticate the signature, Bob generates coherent states of amplitude α with the relative phase defined by the declared private key, and interferes them individually with the states he has in his quantum memory. He monitors the number of photodetection events on his signal null-port arm and confirms the authenticity of the message (*i.e. the message passes authentication*) if the number of photodetection events was below $s_a L$. The parameter s_a is called the **authentication threshold.**

5. To prove to Charlie that he received the message $m = 0$ from Alice, Bob forwards to Charlie the pair $(0, PrivKey_0)$ he received from Alice. Charlie then performs an analogous procedure to Bob, and he verifies the message (*i.e. the message passes verification*) if his number of photodetection events is below $s_v L$ where s_v is called the **verification threshold**, with $0 < s_a < s_v < 1$.

The schematic diagram of the implementation is presented in Figure 7.1 .

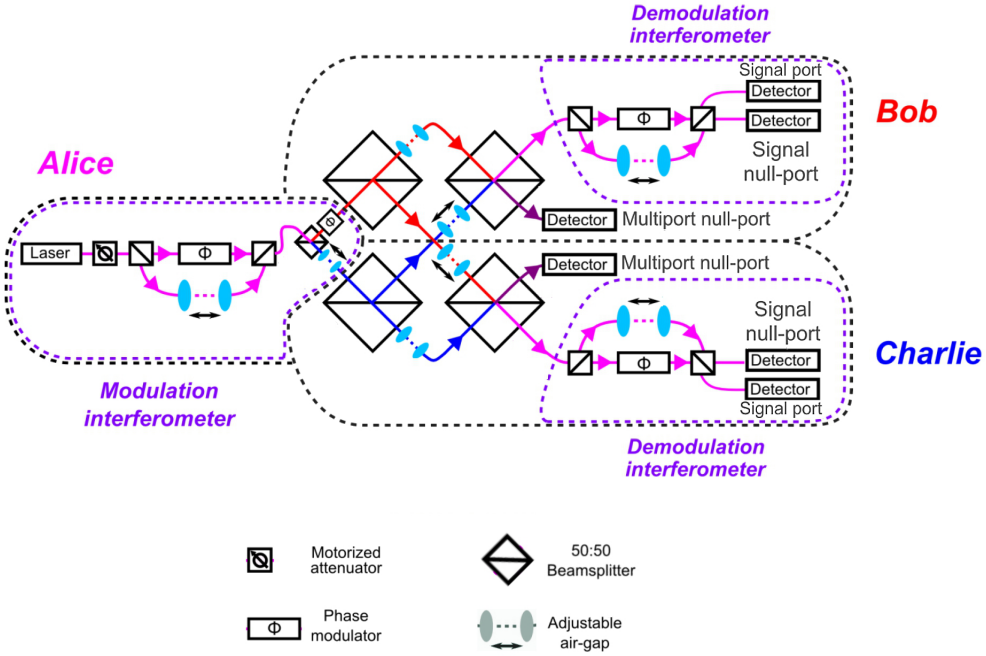

Figure 7.1. The schematic diagram of our experimental set up. Alice's modulation interferometer produces time-interlaced reference and signal pulses. These pulses are split off and sent into the multiport. Bob and Charlie's demodulation interferometers effectively phase modulate the reference pulses and interfere them with the signal pulses. Detection at the signal null-port signifies a phase miss-match. Multiport null-port arms are equipped with photon detectors which can be used to prevent some of the more general forging attacks.

Shared components The system was implemented in polarization-maintaining optical fibre. This is our medium for sending coherent states which comprise the quantum signatures. A common phase reference pulse is realized by time multiplexing reference and signal (phase encoded) pulses and is generated by Alice [106].

Alice's components Alice's (approximations of) coherent states are generated by a vertical cavity surface emitting laser (VCSEL) [107] emitting at a wavelength of 849.8 nm, and with a spectral full-width at half maximum (FWHM) of 0.23 nm. Her emission system operates at a pulse repetition frequency of 100 MHz. The generated pulse is attenuated, in order to achieve the desired mean photon number of the pulses. Alice time multiplexes a phase reference pulse between successive 100 MHz clocked signal pulses using an asymmetric double Mach-Zehnder

approach as used in many quantum key distribution systems employing phase basis sets [106], the set up for which is illustrated in Figure 7.1.

The value $|\alpha|^2$ needs to be defined after Alice's phase modulator/air-gap, as only at this point is any relevant information encoded in the states.

The phase modulated signal pulses, and unmodulated reference pulses are split on a beamsplitter producing the two copies of interlaced quantum signature elements and reference pulses which are forwarded to Bob. One of the output arms of the final beamsplitter is equipped with a phase modulator, and both arms have air-gaps. These were used in some of the experiments to test the performance of the multiport. Additionally, the air-gap in the other arm allows the transmission losses to be balanced between each arm so that the same pulse amplitude is launched to each recipient, and permits compensation for small path-length differences between the two launch arms.

The same electrical signal driving the initial pulse phase modulators (controlled by a pseudo-random number generator) is forwarded to Bob and Charlie. This constitutes the run-time sending of the private key to be authenticated.

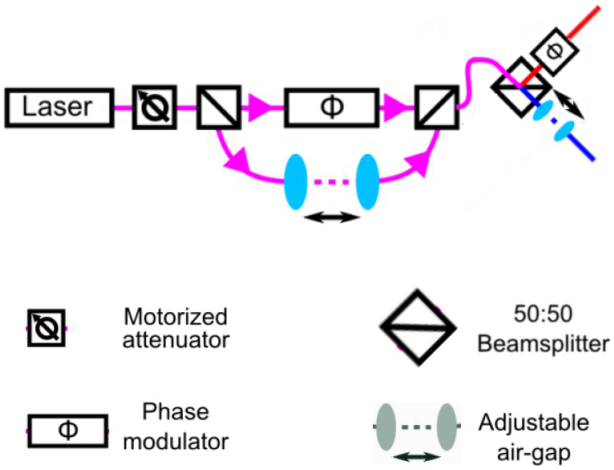

Figure 7.2. Alice's modulation interferometer. An inbound pulse from the laser is split into a signal and a reference pulse. The signal pulse is modulated according to a chosen phase, controlled by a pseudo-random number generator and the pulses are combined time-delayed to prevent symbol interference.

Bob and Chalie's shared multiport The multiport has been assembled out of polarization maintaining optical fibre and balanced 50:50 beam splitter cubes. To ensure a high interferometric fringe visibility in the interferometers comprising the system, it is necessary to ensure that the relative path-length differences remain constant to within a fraction of the emission wavelength of the source laser, and that the light maintains a high degree of polarization [108]. Adjustable air-gaps in active feedback loops are used to compensate for any slow time dependent variations

in the relative path lengths of the interferometers [109]. High extinction ratio linear polarizers are used prior to the motorised optical attenuator to suppress one of the two orthogonal polarization modes emitted by the VCSEL. The fringe visibility is monitored during operation of the system and when a deviation from the expected value is obtained, signature distribution is halted and tuning carried out using a higher intensity signal with a series of known phase modulations until the optimum visibility is obtained.

The two multiport null-port arms are equipped with silicon single-photon avalanche diodes (Si-SPADs) performing the photon detection. Such detectors have previously been used successfully in quantum information experiments [110, 111].

Bob/Charlie's authentification/verification device As we mentioned, the electric signal controlling the modulation of the signal pulses is forwarded to Bob (Charlie) directly. Authentication/verification is performed in run-time using the demodulation interferometer illustrated in 7.1. This interferometer simply inverts the modulation procedure Alice performed, by modulating the reference pulse with the phase which was used to encode the neighbouring signal pulse, and interfering the two pulses on a beamsplitter. The output arms of the final beamsplitter cube are equipped with two Si-SPADs by which mismatches of the phases, and total photonic throughput can be gauged.

As implemented some of the (retrodictively speaking) photons, which take non-interfering paths in sender and receiver (i.e. both short paths or both delayed paths) contribute nothing to the signature (no interference occurs) and these are software gated from the photon arrival times recorded using the free-running SPADs. In post-processing, the time gating software opens a window of duration 2 ns centred on the expected arrival time of a pulse and disregards events which occur outside of this window. As we will mention in Chapter 9, this effectively causes a loss of half of the information-carrying signal. A different implementation of the verification/authentication process may increase the overall efficiency of our system.

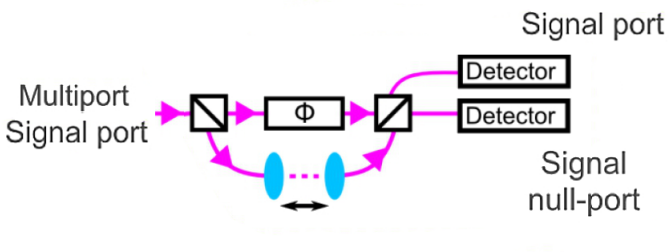

Figure 7.3. Demodulation interferometers. Bob and Charlie's demodulation interferometers are the time-reversal of Alice's modulation interferometer. By precise time-gating, this set-up allows for the detection of phase missmatches of the signal pulse and the reference pulse modulated by a classical control signal.

Signal attenuation in the system The losses of the system have been estimated at 13 dB from the comparison stage input to each demodulation interferometer and 7 dB for each demodulation interferometer. The main contributors to the losses are most likely the air-gaps, used for slow time-dependent variations in the interferometer path lengths and certain experiments, the imperfect splicing of the optical fibres, and the losses specific to the optical components themselves.

For more details on the experimental properties of our system we refer the reader to [6].

7.2 Experiments relevant for the security analysis

For the security analysis of our system, two experimentally measured types of parameters are particularly relevant: the sensitivity of the demodulation interferometers shown in 7.1 to a mismatch in the phases of the compared pulses, and parameters characterising the performance of the multiport itself.

Sensitivity of the demodulation interferometers For the setting of number of distinct phases Alice can generate $N = 8$ and a mean photon number $|\alpha|^2 = 0.16$ we have experimentally counted the number of photon events in Charlie's demodulation interferometer at both the signal port and signal null-port detectors. This was performed for all possible combinations of mismatches and matches between the initial phase modulation caused by Alice, and the demodulation realized in Charlies's demodulation unit.

For each combination of input phases, counts were recorded in a duration of one second. Given the driving frequency of 100 MHz, this yielded the total number of 10^8 generated pulses. By taking the ratio of the number of pulses detected at the null-port arm of the demodulation interferometer and the total number of pulses, the cost matrix in Figure 7.4 was generated. Each entry represents an estimation of the probability of causing a photon detection event per one pulse, given a particular combination of phases of the input pulses, generated over a large sample (10^8). The experimental results presented here are used in 8.3.2.

$\theta' \setminus \theta$	0	$\pi/4$	$\pi/2$	$3\pi/4$	π	$5\pi/4$	$3\pi/2$	$7\pi/4$
0	0.00389	0.0044	0.00524	0.00595	0.00635	0.006	0.00529	0.00439
$\pi/4$	0.00456	0.00388	0.00443	0.00529	0.00604	0.00639	0.00602	0.0052
$\pi/2$	0.00528	0.0046	0.00389	0.00442	0.00529	0.00602	0.00637	0.00595
$3\pi/4$	0.00568	0.00522	0.00458	0.0039	0.0044	0.00524	0.00591	0.0063
π	0.00636	0.00568	0.00527	0.00459	0.00389	0.00443	0.00524	0.00601
$5\pi/4$	0.00562	0.00636	0.00566	0.00523	0.00457	0.00389	0.00441	0.0053
$3\pi/2$	0.00526	0.00568	0.0064	0.0057	0.00522	0.0046	0.00388	0.0044
$7\pi/4$	0.00461	0.00524	0.00565	0.00636	0.00568	0.00522	0.00456	0.00388

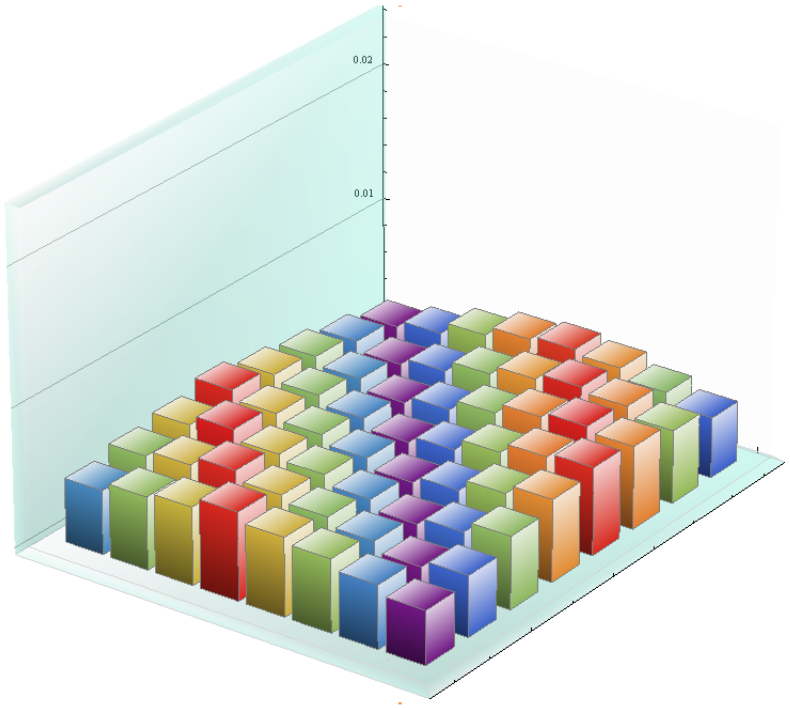

Figure 7.4. The matrix above presents the estimated probability of detecting a photon event at the signal null-port detector, as a function of a phase mismatch between the modulation (rows) and reference demodulation (columns) phases, per individual pulse.

Multiport performance In a second series of experiments, the performance of the multiport itself was investigated, in particular gauging differential losses, which may occur through the multiport going to Bob's or Charlie's null-port or signal-port arms, as illustrated in Figure 7.5. Such asymmetric behaviour, relative to Bob's and Charlie's outputs, can reduce the security properties of the protocol run on this system, as we explain in the security analysis.

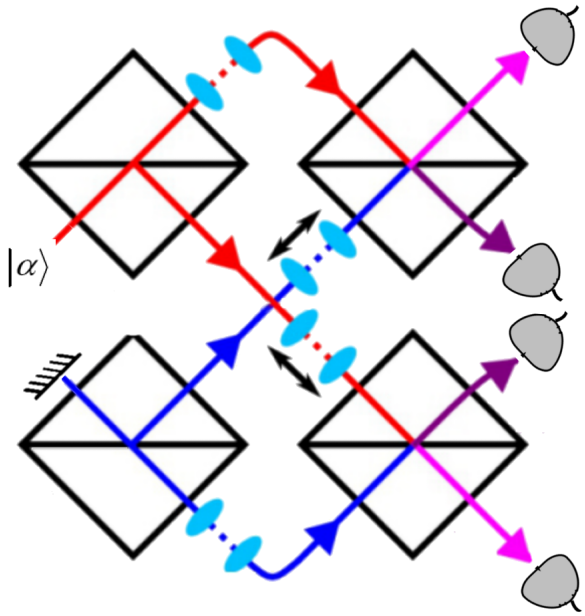

	Charlie's Multiport signal	Bob's Multiport signal	Charlie's Multiport null-port	Bob's Multiport null-port
Air-gap arm blocked	9.1×10^5	2.84×10^5	9.6×10^5	2.61×10^5
Phase modulator arm blocked	1.01×10^6	6.25×10^5	8.5×10^5	7.17×10^5

Figure 7.5. Both input ports were blocked, at different times, the image shows air-gap arm blocked. The table shows a summation of the counts at the 4 detectors.

The results of these two experiments were directly used in the estimations of security against forging (for passive attacks) and security against repudiation and active forging, respectively. The details will be addressed in the following chapter.

Chapter 8

Security analysis of the experiment

We give a detailed analysis of the security properties of a QDS protocol performed on our experimental set-up

8.1 Fundamentals

In this section we provide the security analysis of the three party QDS protocol realized using coherent states. Concretely, we calculate the probabilities of successful forging, and repudiation by malevolent parties as functions of private key lengths, and show that the presented protocol is asymptotically robust. The security analysis assuming passive attacks, defined later in this chapter, is based on observed experimental results. The preliminary analysis of security under general types of attacks is performed under certain assumptions concerning noise models in our system, consistent with the experimental observations.

8.1.1 Protocol outline

The outline of the protocol is given in Protocol 14.

8.1.2 Definitions of security

The presented Quantum Digital Signatures protocol is designed to be immune to two types of malicious activities: forging and repudiation. Immunity to forging signifies that any receiving party will reject any message which was not sent by honest Alice herself. Immunity to repudiation signifies that if Alice sends a message to Bob which passes authentication, afterwards the same message will pass verification with Charlie as well. In other words, Alice cannot make Bob and Charlie disagree on the authenticity, and consequently the content, of her message. Naturally, the protocol has to also be robust, meaning, if all parties play honestly, no message is rejected even in the presence of imperfections. More formally we have the following:

Protocol 14 Quantum Digital Signatures with Phase-Encoded Coherent States

1. To sign a single bit (message $m = 0$ or 1) in the future, Alice generates two sequences $PrivKey_0 = (\theta_1^0, \ldots \theta_L^0)$ and $PrivKey_1 = (\theta_1^1, \ldots \theta_L^1)$ of L randomly chosen angles from the set of N equally spaced phases, so

$$\theta_k^m \in \left\{ \frac{2r\pi}{N} | r = 0, \ldots, N-1 \right\}. \tag{8.1}$$

   The pair $(m, PrivKey_m)$ is called a private key pair for message m.

2. Alice then generates two copies of a sequence of coherent states $QuantSig_0 = (\rho_1^0, \ldots, \rho_L^0)$ with the coherent phases matching the angles in the sequence $PrivKey_0$, thus $\rho_k^0 = |e^{i\theta_k^0}\alpha\rangle\langle e^{i\theta_k^0}\alpha|$ where α is a real positive amplitude. A sequence of such states is called a quantum signature. She sends a copy of the quantum signature to each of Bob and Charlie each, informing them that they correspond to message $m = 0$. Alice then does analogously for the message $m = 1$. The individual state ρ_k^m we refer to as the k^{th} quantum signature element state for message m.

3. Bob and Charlie send their copies of the sequences $QuantSig_0$ and $QuantSig_1$ through the multiport, saving the states exiting the multiport signal arm in quantum memory, noting which quantum signature corresponds to message $m = 0$ and which to $m = 1$. The multiport null-ports on Bob's and Charlie's side are equipped with photon detectors and the total number of photon events here will serve to disable certain types of forging attacks, but are not crucial for security against message repudiation. For the simple case of passive attacks which we define and analyze first, these outcomes will be ignored.

4. To sign a single bit, say $m = 0$ with Bob, Alice announces the message m and the corresponding private key to Bob (thus she sends the pair $(0, PrivKey_0)$ over an untrusted channel). To authenticate the signature, Bob generates coherent states of amplitude α with the relative phase defined by the declared private key, and interferes them individually with the states he has in his quantum memory. He monitors the number of photodetection events on his signal null-port arm and confirms the authenticity of the message (***i.e. the message passes authentication***) if the number of photodetection events was below $s_a L$. The parameter s_a is called the **authentication threshold.**

5. To prove to Charlie that he received the message $m = 0$ from Alice, Bob forwards to Charlie the pair $(0, PrivKey_0)$ he received from Alice. Charlie then performs an analogous procedure to Bob, and he verifies the message (***i.e. the message passes verification***) if his number of photodetection events is below $s_v L$ where s_v is called the **verification threshold**, with $0 < s_a < s_v < 1$.

If any of the thresholds are breached, the protocol is aborted.

- We say that a protocol realizing QDS is secure against forging if the probability of a recipient successfully producing, without receiving it from Alice, a private key of message m which will pass verification by the other recipients is *decaying exponentially quickly in terms of the quantum signature length* L .
- We say that a protocol realizing QDS is robust if in the setting where all parties are honest, a message will be authenticated and verified *except with probability decaying exponentially quickly in terms of the quantum signature length* L .

For simplicity we will always consider Bob to be the forger. Note that any security can only be guaranteed if only one party is cheating - two cooperating parties can always cheat on the third. Thus, when analysing security against forging, Alice is assumed to be honest, and in the security against repudiation, Bob and Charlie are assumed to be honest.

Assumptions on quantum and classical channels Throughout the analysis of this chapter, we assume that the quantum channel from Alice to an individual recipient is under the recipients control during the distribution step (Step 2 in Protocol 14).

More precisely, while the quantum channel is not assumed to be private, Alice and the recipient have some means to ensure that an external party is not tampering with the states sent over this channel. This assumption is crucial: if one recipient has the power to interfere with the quantum channels leading from Alice to other recipients, he can perform a "key-swap attack", and substitute Alice's quantum signatures with his own. This would constitute a complete breaking of the security. How and if such an authentication scheme could be realized is addressed later in Section 9.3.

As mentioned in the introduction, this type of an assumption is standard in public-key cryptography, and there, without the assumption that the public keys have been distributed without tampering, a key swap attacks becomes possible. However, as we have noted, since the "public-keys" in QDS comprise quantum states it is not obvious whether such an assumption can or cannot be justified. These types of considerations, while crucial, go beyond the scope of research presented in this thesis. Nonetheless, the question of how quantum signature authentication could be performed we will address briefly in the Chapter 9. Additionally, it is assumed Alice cannot tamper with the classical channels used in the verification step (Step 5 in Protocol 14). This can be ensured by having these messages authenticated using QDS itself, by having every party in the group establish quantum signatures with every other party. Thus, the fundamental assumptions we work with are analogous to the standard assumptions in public-key cryptography (but admittedly not the same). The cost and consequences of the last two assumptions without any assumptions, but rather in terms of the cost of quantum message authentication schemes employing pre-shared private keys are addressed in Section 9.3.

8.2 Cheating Alice – security against repudiation

The formal definition of security against message repudiation is given in terms of a conditional statement: if one recipient party authenticates the message (say Bob), the other party (Charlie) will verify it as well. This definition agrees with the initial security requirements given in [2, 3] and constitutes the security guarantee to the recipients in the protocol. Thus, successful repudiation means that she sends a message, say $(0, PrivKey_0)$, to Bob, he checks it and forwards it to Charlie, who then rejects. For the remainder of this section we analyze the probability of this happening. The robustness of the protocol, that is abort probability in the honest setting caused by imperfections, shall be addressed later on in this chapter.

The most general state Alice can prepare is

$$\pi_{A,B_1,C_1,\ldots B_L,C_L}, \tag{8.2}$$

which is a general $2L+1$-partite state. Subsystem A Alice keeps, and sends partitions $B_1, \ldots B_L$ to Bob and $C_1, \ldots, C_L$ to Charlie. To clarify, for instance if Alice is honest there is no subsystem A, and C_i and B_i are identical coherent states with a complex phase known to Alice alone, as specified by the protocol.

8.2.1 Security against repudiation – separable attacks

We first assume that the multiport Bob and Charlie have is ideal and that the system A is disentangled from the rest of Alice's state (or simply does not exist), and the subsystems (B_kC_k) and (B_lC_l) are not entangled with each other for $k \neq l$.

However, we allow the partitions B_k and C_k to be mutually entangled. This type of an attack we refer to as a separable attack. According to the protocol specifications Charlie and Bob will individually run the pairs of states in the systems (B_kC_k) through the multiport, and commit to quantum memory whatever comes out at their signal outputs of the multiport. For the purposes of showing security against repudiation, we can assume that they ignore the measurement outcomes on the multiport null-ports.

For the k^{th} signature element, the joint system of Charlie and Bob which they store into memory is some state $\pi^{out}_{B_k,C_k}$ which is symmetric under permutations of Bob's and Charlie's subsystems, as we now show. Let

$$\pi^{in}_{B_kC_k} = \int_{\mathbb{C}^2} P(\alpha,\beta)|\alpha\rangle\langle\alpha| \otimes |\beta\rangle\langle\beta| d^2\alpha d^2\beta$$

be any general two mode state given in the P representation, which enters the multiport. Then the stored output state is (*i.e.* when the null-port subsystems have been traced out):

$$\pi^{out}_{B_kC_k} = \int_{\mathbb{C}^2} P(\alpha,\beta)|(\alpha+\beta)/\sqrt{2}\rangle\langle(\alpha+\beta)/\sqrt{2}| \otimes |(\alpha+\beta)/\sqrt{2}\rangle\langle(\alpha+\beta)/\sqrt{2}| d^2\alpha d^2\beta \tag{8.3}$$

which is symmetric in the sense given above. In the process of checking Alice's message, Bob and Charlie will perform a sequence of measurements on their subsystems, and the measurements will be identical as they are prescribed by the (same) private key Alice had sent. Since the systems Bob and Charlie have are symmetric under the swap of their subsystems, the probability matrix of their respective outcomes will be symmetric as well. To explain this let us focus on the k^th subsystem $\pi^{out}_{B_kC_k}$ given in 8.3. Let, say Bob be the first to check this subsystem, by preparing some coherent state prescribed by the (common) private key Alice gave them, and interfere that state with his corresponding subsystem of the public key, namely $Tr_{C_k}(\pi^{out}_{B_kC_k})$, checking whether he gets a click on his signal null-port. This constitutes a measurement, characterised by some two-outcome POVM Π_0, Π_1 corresponding to the setting of getting a click or not, applied on the state $Tr_{C_k}(\pi^{out}_{B_kC_k})$. If Charlie was the first to check, he would have done the same measurement on the subsystem $Tr_{B_k}(\pi^{out}_{B_kC_k})$. But, since $\pi^{out}_{B_kC_k}$ is symmetric under subsystem swap, the probability matrix of the join four outcome measurement $\Pi_{i,j} = \Pi_i \otimes \Pi_j$ for $i, j = 0, 1$, is symmetric as well, so for every possible state $\pi^{out}_{B_kC_k}$ the probability of getting outcomes $(0, 1)$ (corresponding to the POVM element $\Pi_0 \otimes \Pi_1$) and $(1, 0)$ (corresponding to the POVM element $\Pi_1 \otimes \Pi_0$)is the same.

Assume Alice wishes Bob to accept and Charlie to reject. Alice requires Charlie to accumulate more photodetection events than Bob. Then the a-priori probability of Bob not detecting any photons, and Charlie detecting one or more photons, is no higher than $1/2$ (as the opposite event must be equally likely), per each pair of pulses. According to the protocol specifications, Bob accepts if he gets less than s_aL photodetection events.

Charlie accepts at less than s_vL photodetection events. Thus, Charlie needs to accumulate $(s_v - s_a)L$ photodetection events more than Bob in order for Alice's cheating to succeed. The choice of values of s_a and s_v come from the security analysis against forging and will be calculated later. The probability of Alice achieving her goal of getting, say Bob to accept, and Charlie to reject is then $d^{(s_v-s_a)L}$. where d is the probability of getting the outcome $(0, 1)$ and, as we have shown, the outcome $(1, 0)$ as well. The value

$$d^{(s_v-s_a)L} \tag{8.4}$$

is maximized for the highest allowable value of d, so for $d = 1/2$, yielding an overall probability of Alice cheating successfully as [1]

$$\epsilon_{repudiate} = \left(\frac{1}{2}\right)^{(s_v-s_a)L}. \tag{8.5}$$

8.2.2 Security against repudiation – coherent attacks

In a coherent attack, the entanglement of the states Alice may use is unrestricted. The probability of Alice cheating, in the case of restricted entanglement of the states as previously described,

[1] The formula below holds exactly when $s_a = 0$ and approximately when s_a is small, as is in our case. An exact formula for the bound is not expressible in a closed form, and a weak upper bound (for large s_v) attains the form similar to the forging probability we will express shortly.

does not change if we slightly modify the protocol:

1. Alice first sends the elements of the quantum signatures,
2. Bob and Charlie run them through the multiport,
3. Alice sends the corresponding angle to both Charlie and Bob and they immediately measure,
4. and at each step, for each element of the quantum signature, Alice learns the outcomes of Bob's and Charlie's measurements.

This modification does not increase a malevolent Alice's cheating probability. Thus we have

$$P(Alice\ Cheats|original\ protocol,\ separable\ attack) = \tag{8.6}$$
$$P(Alice\ Cheats|modified\ protocol,\ separable\ attack), \tag{8.7}$$

and we will next prove

$$P(Alice\ Cheats|modified\ protocol,\ separable\ attack) = \tag{8.8}$$
$$P(Alice\ Cheats|modified\ protocol,\ coherent\ attack). \tag{8.9}$$

Finally, we will show that the modified protocol can only help Alice in the coherent attack, leading to

$$P(Alice\ Cheats|modified\ protocol,\ coherent\ attack) \geq \tag{8.10}$$
$$P(Alice\ Cheats|original\ protocol,\ coherent\ attack). \tag{8.11}$$

This proves that coherent attacks cannot help Alice.

As noted, the most general state Alice could use in her attempt to cheat is $\pi_{A,B_1,C_1,\ldots B_L,C_L}$. The subsystem A remains with Alice, and the rest is sent to Bob and Charlie and will traverse the multiport. First of all, note that in the original protocol, there is no interaction between Alice on one side and Bob and Charlie on the other, once she has declared her private key. If Alice had a system A which is still entangled with whatever Bob and Charlie save after the multiport action, the action of measurement by Bob and Charlie in the verification part cannot convey any information to Alice through the system A, since she does not learn the outcomes of Bob's and Charlie's measurements. Hence, she cannot gain anything by manipulating system A, and in the original protocol we may assume that Alice simply send the state

$$Tr_A\left[\pi_{A,B_1,C_1,\ldots B_L,C_L}\right].$$

We will now show that if she uses a separable strategy in the modified protocol, Alice can achieve the same measurement statistics during verification and authentication as by using a coherent attack. Initially, we assume an ideal multiport. The first state Bob and Charlie may measure is $Tr_{B_2,C_2,\ldots B_L,C_L}\left[M(\pi_{B_1,C_1,\ldots B_L,C_L})\right]$ where M denotes the global action of the multiport. Alice

could simply have sent this state to Charlie and Bob and achieved the same measurement statistics as for this state (any state she sends which is already symmetric will not be changed by the multiport). However, the measurement outcome may influence the rest of the system, which Alice has not yet sent to Bob and Charlie. But, if the authentication and verification measurement outcomes are revealed to Alice at each step, then the state of the rest of her system is also known to her at each step. Then in the sequential setting, she can prepare the corresponding signature state for the second measurement and attain the same measurement statistics. This continues inductively.

Thus, we have shown the following: any measurement statistics achieved using a globally entangled cheating state can be achieved using a separable attack, if Alice is allowed to learn the measurement outcomes before sending the next pair of states. This proves the required claim

$$P(Alice\,Cheats\,|modified\,protocol,\ individual\ attack) = \tag{8.12}$$

$$P(Alice\,Cheats|modified\,protocol,\ coherent\ attack). \tag{8.13}$$

To finalize our proof we need to show that

$$P(Alice\,Cheats|modified\,protocol, coherent\ attack) \geq \tag{8.14}$$

$$P(Alice\,Cheats|original\ protocol, coherent\ attack). \tag{8.15}$$

This is easy to see as by simply ignoring the information Alice additionally gets in the modified protocol, what Alice runs is effectively the original protocol, barring the timing of the measurements. However, the timing cannot influence the measurement statistics, and hence cannot influence Alice's cheating probability. So our claim holds and using globally entangled states cannot help Alice repudiate her signed messages.

To summarize, as long as the properties of robustness and security against forging can be maintained for some s_v strictly greater than s_a, then security against repudiation can be guaranteed as well.

8.2.3 Security against repudiation with realistic devices

In the security analysis against repudiation, for the ideal, case the crux of the argument is that the states Bob and Charlie share are symmetric under swap of their respective subsystems. This guaranteed that for a single pair of states, the probabilities of the outcomes $(0, 1)$ and $(1, 0)$ of joint measurements made by Bob and Charlie are equal. Since these are equal, each is at most $1/2$, and this value is raised to the exponent $(s_v - s_a)L$ to obtain the upper bound on the probability of Alice successfully cheating (for details see Section 8.2.1). Here, we briefly address the effects imperfect realization may have on the security of our system. Note that the multiport acts as a CPTP map (completely positive trace preserving map, quantum channel) on the input state, where the output state is the joint state of the Bob's and Charlie's multiport signal outputs, *i.e.* elements of the quantum signature.

Let M_{ideal} and M_{real} be the corresponding CPTP maps of the ideal and real multiport.

For any input state π_{in} we have that $M_{ideal}(\pi_{in})$ realizes a symmetric probability matrix with respect to outcomes of identical measurements done by Charlie and Bob. Assume that Alice wishes Bob to not register a photodetection event while Charlie does. This probability for the state $M_{ideal}(\pi_{in})$ is at most $1/2$. Let the probability of the same event for the state $M_{real}(\pi_{in})$ be d. In this case, the probability of Alice cheating is $d^{(s_v - s_a)L}$. By the properties of the trace distance we have that $D(M_{real}(\pi_{in}), M_{ideal}(\pi_{in})) \geq |1/2 - d|$, where $D(\rho, \eta)$ denotes the trace distance between the states ρ, η. For details on the notion of the trace distance and the derivation of this claim see Section 8.4.2.

Thus, as long as $D(M_{real}(\pi_{in}), M_{ideal}(\pi_{in})) < 1/2$ the probability of Alice cheating will diminish exponentially quickly in terms of L. One way to conclusively show that $D(M_{real}(in), M_{ideal}(in)) < 1/2$ holds for our system would be to use full process tomography, which was not performed. (Full process tomography for CV systems is not as well investigated as for qubits, and even qubit process tomography is experimentally demanding.) To evaluate where the actual worst case value d may lie for our implementation, we instead analyse how different types of possible imperfections influence this parameter. The imperfections may in general occur within the multiport, but also during the processes leading to Bob and Charlie finally detecting or failing to detect photons, that is, the events of interest $(0, 1)$ and $(1, 0)$.

In principle M_{real} can be written as a composition of M_{ideal} with noise/loss CPTP map collecting all the effects caused by the imperfections in our system. The imperfections characterising the noise/loss map are brought about by the imperfections within the multiport itself. Additionally we also consider the imperfections caused by the realistic authentication/verification process and their effect on the security against repudiation. If identical sets of equipment are used in Charlie and Bob for the purposes of authentication/verification then the losses and noise induced in the individual arms act as identical and uncorrelated (separable) CPTP maps on the states exiting the multiport. This process can only reduce the trace distance between the reduced states of Bob's and Charlie's systems, thus such noise can only reduce a malevolent Alice's success probabilities. For the purposes of upper bounding the repudiation probability (*i.e. worst-case scenario*) we may ignore uncorrelated imperfections associated with authentication and verification, and the only imperfections which may help Alice have to lie within the multiport itself. Here, again, any imperfection causing identical uncorrelated noise/loss cannot help Alice, by the same arguments as above. Hence, we only need to focus on correlated, or differential imperfections inducing correlated or unequal CPTP maps contributing to the cumulative noise/loss map on Bob's and Charlie's reduced states. In our implementation of the multiport, the most likely culprit of differential imperfections comes from the variable air gaps and attenuators placed into the arms of the interferometers. The optical attenuators compensate for different loses in the optical components ensuring the equal intensity of interfering beams. The air gaps compensate for variations in the optical path length in the interferometers which arise from environmental fluctuations. These technical necessities primarily induce an uneven loss in both signal and reference pulses with Bob and Charlie respectively, and this is the effect we now focus on. This differential loss was studied by the experiment explained in Figure 7.5 and Section 7.2, paragraph Multiport

performance.

We can see in Table 7.5 that differential loss causes Bob to receive on average no less than $1/4$ of the photons compared to Charlie. Since both the signal and the reference pulse are identically attenuated this can, in the worst case scenario, cause the event $(0, 1)$ to be ten times more likely than $(1, 0)$. If Alice wishes to repudiate her message with the party with the lower loss (Charlie), this induces the worst case value of $d = 4/5$. Even if Bob's and Charlie's output losses were a thousand-fold different (inducing the value $d = 1000/1001$), the forging probability as a function of the signature length L is significantly higher than the refutation probability, which will become clear from the computations to follow. Thus if one is interested in probability of the protocol failing in any way, security against forging, and likewise the required robustness, will constitute the dominant factor in the overall failure probability of the protocol. Forging is therefore the focus of the remainder of this chapter.

8.3 Cheating Bob – security against forgery

We identify two types of cheating strategies for forger Bob:

> **Passive strategy:** Bob does not interfere during distribution of the quantum signatures, but tries to cheat by inspecting his copy of the quantum signature. Also, they capture Bob's cheating capabilities if the quantum signatures were distributed to the recipients vis a trusted centre, without the use of the multiport.
>
> **Active strategy:** Bob is malevolent throughout the distribution of the quantum signatures - this constitutes the most general type of attacks.

In the first type of strategies for Bob, we further identify separable and collective types of attacks, collective being more general. We show that collective attacks in passive strategies do not help.

For active strategies, we will distinguish between separable and coherent strategies, the latter being completely general. Bob can benefit from active attacks relative to passive strategies. In this type of attacks, we offer a security proof for separable strategies, and for coherent strategies we give with a plausibility argument only, that they cannot help Bob. The separable vs. collective/coherent strategies here are somewhat analogous to the strategies a malevolent eavesdropper Eve may resort to in QKD [112].

We begin with analysis of the passive attack, the results of which will be the crux of the security analysis for active attacks.

8.3.1 Passive strategy - separable attacks

In this type of an attack, Bob does not interfere throughout the quantum signature distribution phase, so he lets the states he receives traverse the multiport unperturbed. To forge a message, he applies one (optimal) measurement to estimate the phase of each of his elements of the quantum

signature and sends his best guess to Charlie. Thus, to calculate Bob's cheating probability for this attack we need to calculate the probability of Bob not generating a photodetection event with Charlie, per individual quantum signature element. This probability is given by

$$p_{forgery} = \min_{\{\Pi_\phi\}} \frac{1}{N} \sum_{\phi} \sum_{\theta} Tr(\Pi_\phi \rho^\theta) c_{\phi,\theta} \tag{8.16}$$

where

- $Tr(\Pi_\phi \rho^\theta)$ is the probability of Bob measuring (and thus declaring) the angle ϕ if the state he measured was in fact encoded with the angle θ,
- $c_{\phi,\theta}$ is the probability of Charlie getting a photon detection event in his signal null-port if the state he had in his memory was encoded with θ and Bob sent ϕ.

The expression above is minimized over all possible POVMs.

The value $p_{forgery}$ constitutes the cost of a *minimum cost measurement*, and the criteria on the POVM elements (the measurements) for which the minimum is achieved are given in [113] with

1. $\Gamma = \sum_i \Pi_i W_i = \sum_i W_i \Pi_i$ for $W_i = \sum_j C_{i,j} \rho_j$
2. $\Gamma = \Gamma^\dagger$
3. $\Pi_i (W_i - \Gamma) = (W_i - \Gamma)\Pi_i = 0 \; for \; all \; i$, and
4. $(W_i - \Gamma)$ is positive-semidefinite for all i

We refer to the set of conditions above as Helstrom criteria 1-4, respectively. The cost matrix $C = [c_{\phi,\theta}]_{\phi,\theta}$ is obtained from experimental results. For clarity, the cost matrix is indexed according to encoding angles. However, to remain compatible with the indexing tradition for minimum cost measurements, in the abstract formulation of the problem the indexing is performed across integers, so that the index angle $\theta = 2k\pi/N$ corresponds to the integer index k.

In the most general case for an arbitrary cost matrix, the computation of the optimal measurement is difficult. However, note that if the cost matrix C is replaced by a cost matrix, where each entry is less than or equal to the entries of the original cost matrix (an element-wise dominated matrix), the overall cost of the optimal transform can only decrease.

In the ideal case, where the experiment is completely symmetric, the cost matrix C is circulant and symmetric. A circulant matrix is a square matrix whose each row is a cyclic right-shift of the previous row. For the formal definition of a circulant matrix, and useful mathematical properties it has, see Section 10.

However, in reality the matrix C it is just close to a symmetric and circulant matrix. If we now substitute the cost matrix with the closest element-wise dominated symmetric and circulant matrix, and compute the cost for this matrix, by the remark above we have found a lower bound for $p_{forgery}$. In a similar fashion, we can compute the upper bound for the same expression by considering the symmetric and circulant cost matrix which upper bounds the elements of the

actual cost matrix C. As we will see, these two values are very close, so the lower bound we will compute is very close to the actual value.

This reduction simplifies the computation of the bounds on the cheating probability as for circulant and symmetric positive cost matrices the first, second and third Helstrom criteria are satisfied for the so-called minimum-error or square-root measurement [113]. This claim is proven in Section 8.4.3, where the square-root measurement is also given.

The conditions on the cost matrix for the fourth Helstrom criterion to be satisfied with the square-root measurement are more involved, and for our experimental results we have verified this criterion numerically [2].

If Bob were honest, the probability of him triggering a photodetection event with Charlie, per quantum signature element state, in Charlie's signal null-port, is given by the average of the diagonal of the cost matrix C. Let this value be $p_{original}$, and let the corresponding value, if Bob is forging, be $p_{forgery}$.

We define the gap between these two values: $g = p_{forgery} - p_{original}$. If we now set the authentication and verification thresholds at

$$s_a = p_{original} + 1/3g \tag{8.17}$$

and

$$s_v = p_{forgery} - 1/3g = p_{original} + 2/3g, \tag{8.18}$$

the probability of Bob successfully forging the signature is equal to the probability that the fraction of photon detection events is less than s_v where the expected fraction is $p_{forgery}$. This is then the probability that in a repeated experiment (L times) with a binary outcome with mean $p_{forgery}$ the normalized measured outcome diverges from the expectancy by more than $p_{forgery} - s_v = 1/3g$ and this is bound using the Hoffendig's bound as follows:

$$\epsilon_{forging} = P_{Bob\ cheats} \leq 2\exp(-\frac{2}{9}g^2 L). \tag{8.19}$$

A similar analysis gives us the robustness as well:

$$\epsilon_{robustness} = P(Honest\ setting\ abort) \leq \exp(-\frac{2}{9}g^2 L) + \exp(-\frac{4}{9}g^2 L) \tag{8.20}$$

which is bounded above by $\epsilon_{forging}$.

[2] Certain relatively involved analytical conditions can be derived for testing the fourth Helstom's criterion for symmetric and circulant cost matrices which depend on the spectrum of the cost matrix. Since the spectrum can in the end only be computed numerically given that the cost matrix we have comes from experimental data, it was more convenient to check this last constraint directly.

8.3.2 *Estimation of forging probabilities for the passive attack based on experimental data*

The cost matrix realized by our experimental set-up using 8 differing phase states and with average photon number of $|\alpha|^2 = 0.16$ is given with

$$C = \begin{pmatrix} 0.00389 & 0.0044 & 0.00524 & 0.00595 & 0.00635 & 0.006 & 0.00529 & 0.00439 \\ 0.00456 & 0.00388 & 0.00443 & 0.00529 & 0.00604 & 0.00639 & 0.00602 & 0.0052 \\ 0.00528 & 0.0046 & 0.00389 & 0.00442 & 0.00529 & 0.00602 & 0.00637 & 0.00595 \\ 0.00568 & 0.00522 & 0.00458 & 0.0039 & 0.0044 & 0.00524 & 0.00591 & 0.0063 \\ 0.00636 & 0.00568 & 0.00527 & 0.00459 & 0.00389 & 0.00443 & 0.00524 & 0.00601 \\ 0.00562 & 0.00636 & 0.00566 & 0.00523 & 0.00457 & 0.00389 & 0.00441 & 0.0053 \\ 0.00526 & 0.00568 & 0.0064 & 0.0057 & 0.00522 & 0.0046 & 0.00388 & 0.0044 \\ 0.00461 & 0.00524 & 0.00565 & 0.00636 & 0.00568 & 0.00522 & 0.00456 & 0.00388 \end{pmatrix} \tag{8.21}$$

The symmetrized and circularized cost matrix which lower bounds the original cost matrix is characterised by its first row, given with

$$C'_{row} = (0.00388, 0.00439, 0.0052, 0.00591, 0.0063, 0.00591, 0.0052, 0.00439). \tag{8.22}$$

and the upper bounding symmetrized and circularized matrix is characterised by the row:

$$C''_{row} = (0.0039, 0.00443, 0.0053, 0.00604, 0.00639, 0.00604, 0.0053, 0.00443). \tag{8.23}$$

For both lower and upper bounding cost matrices we have the fourth Helstrom criterion satisfied, so in both cases, the minimum cost measurement is realized by the minimum error measurement POVM's, and the costs are given as follows: $cost_{lower} = 4.7{\times}10^{-3}$ and $cost_{upper} = 4.76{\times}10^{-3}$. As noted, for the worst case scenario, we need to take the largest diagonal element of the actual cost matrix as p_{honest} so it is 3.9×10^{-3} and we have the lower and upper bounds on the gap g as follows: $g_{lower} = 8.03 \times 10^{-4} \pm 2.8 \times 10^{-5}$, $g_{upper} = 8.64 \times 10^{-4} \pm 5.5 \times 10^{-5}$, and this demonstrates that the bounding technique yields a reasonable bound. Thus, the security of our system is characterised by the lower bound on the gap

$$g := g_{lower} = 8.03 \times 10^{-4} \pm 2.8 \times 10^{-5}. \tag{8.24}$$

In the cost matrix above we have just represented the raw data obtained from the experiment. However in the actual calculation of the gap we have included the computed confidence intervals to obtain a final error estimate.

8.3.3 *Passive strategies with collective measurements*

In the security analysis for the passive attack above, we have assumed that the malevolent Bob performs individual identical measurements on his quantum signature states in order to produce a "best guess" sequence of phase angles to use when forging a message. A collective measurement

may in principle yield a higher probability of forging a message, but here we prove this is not the case. This is not a surprising result as the quantum signature element states are not mutually correlated. Recall, the pivotal value, which we used to characterise the security of our system was $p_{forgery}$–the probability of a cheating Bob not causing a photodetection event during Charlie's verification phase, per individual quantum signature element state. We now show that any average probability of a cheating Bob not causing a photodetection event during Charlie's verification phase, per individual quantum signature state, if Bob uses a global measurement, can be achieved by measurements of individual signature states. This shows that collective measurement strategies cannot help a malevolent Bob.

Let $\{\Pi_{\overrightarrow{\phi}}\}$ for $\overrightarrow{\phi} = (\phi_1, \ldots, \phi_L)$ be the POVM elements of any global measurements Bob may employ. where the index is a sequence of angles corresponding to Bob's estimate of the angles. Then the average probability of Bob not causing a proton detection event with Charlie is:

$$p_{forgery}^{average} = \frac{1}{N^L} \sum_{\overrightarrow{\phi}} \sum_{\overrightarrow{\theta}} Tr(\Pi_{\overrightarrow{\phi}} \rho^{\overrightarrow{\theta}}) c_{\overrightarrow{\phi},\overrightarrow{\theta}} \tag{8.25}$$

with $\rho^{\overrightarrow{\theta}} = \otimes_{k=1}^{L} |e^{i\theta_k}\alpha\rangle\langle e^{i\theta_k}\alpha|$ and $c_{\overrightarrow{\phi},\overrightarrow{\theta}} = \sum_{k=1}^{L} c_{\phi_k,\theta_k}/L$. Then we have the following derivation:

$$p_{forgery}^{average} = \frac{1}{N^L} \sum_{\overrightarrow{\phi}} \sum_{\overrightarrow{\theta}} Tr(\Pi_{\overrightarrow{\phi}} \rho^{\overrightarrow{\theta}}) c_{\overrightarrow{\phi},\overrightarrow{\theta}} =$$

$$\frac{1}{N^L}\frac{1}{L} \sum_{k=1}^{L} \sum_{\phi_k} \sum_{\theta_k} \sum_{(\phi_1,\ldots,\phi_{k-1},\phi_{k+1},\ldots,\phi_L)} \sum_{(\theta_1,\ldots,\theta_{k-1},\theta_{k+1},\ldots,\theta_L)} Tr(\Pi_{\overrightarrow{\phi}} \rho^{\overrightarrow{\theta}}) c_{\phi_k,\theta_k} =$$

$$\frac{1}{N^L}\frac{1}{L} \sum_{k=1}^{L} \sum_{\phi_k} \sum_{\theta_k} Tr \left(\sum_{(\phi_1,\ldots,\phi_{k-1},\phi_{k+1},\ldots,\phi_L)} \sum_{(\theta_1,\ldots,\theta_{k-1},\theta_{k+1},\ldots,\theta_L)} \Pi_{\overrightarrow{\phi}} \rho^{\overrightarrow{\theta}} \right) c_{\phi_k,\theta_k} =$$

$$\frac{1}{N}\frac{1}{L} \sum_{k=1}^{L} \sum_{\phi_k} \sum_{\theta_k} Tr \left(\left(\sum_{(\phi_1,\ldots,\phi_{k-1},\phi_{k+1},\ldots,\phi_L)} \Pi_{\overrightarrow{\phi}} \right) \left(\sum_{(\theta_1,\ldots,\theta_{k-1},\theta_{k+1},\ldots,\theta_L)} \frac{1}{N^{L-1}} \rho^{\overrightarrow{\theta}} \right) \right) c_{\phi_k,\theta_k}$$

Note that the operator

$$\tilde{\Pi}_{\phi_k}^{k} = \left(\sum_{(\phi_1,\ldots,\phi_{k-1},\phi_{k+1},\ldots,\phi_L)} \Pi_{\overrightarrow{\phi}} \right) \tag{8.26}$$

is a positive operator, and that

$$\left(\sum_{(\theta_1,\ldots,\theta_{k-1},\theta_{k+1},\ldots,\theta_L)} \frac{1}{N^{L-1}} \rho^{\overrightarrow{\theta}} \right) = \Phi^{\otimes(k-1)} \otimes \rho^{\theta_k} \otimes \Phi^{\otimes(N-k)} \tag{8.27}$$

with $\Phi = 1/N \sum_{\theta} \rho^{\theta}$ being the average quantum signature element state. Thus we have:

$$p_{forgery}^{average} = \frac{1}{N}\frac{1}{L}\sum_{k=1}^{L}\sum_{\phi_k}\sum_{\theta_k} Tr\left(\tilde{\Pi}_{\phi_k}^{k}\Phi^{\otimes(k-1)} \otimes \rho^{\theta_k} \otimes \Phi^{\otimes(N-k)}\right) c_{\phi_k,\theta_k} =$$

$$\frac{1}{N}\frac{1}{L}\sum_{k=1}^{L}\sum_{\phi_k}\sum_{\theta_k} Tr\left(\tilde{\Pi}_{\phi_k}^{k}\left(\Phi^{\otimes(k-1)} \otimes \mathbb{1} \otimes \Phi^{\otimes(N-k)}\right)\left(\mathbb{1}^{\otimes(k-1)} \otimes \rho^{\theta_k} \otimes \mathbb{1}^{\otimes(N-k)}\right)\right) c_{\phi_k,\theta_k} 2$$

where $\mathbb{1}$ is the identity operator. The trace superoperator above can be decomposed into the partial trace over the k^{th} subsystem and the partial trace over every subsystem except the k^{th} subsystem, which we will denote $Tr_{\bar{k}}$:

$$\begin{aligned} p_{forgery}^{average} = \frac{1}{N}\frac{1}{L}\sum_{k=1}^{L}\sum_{\phi_k}\sum_{\theta_k} Tr\left(\tilde{\Pi}_{\phi_k}^{k}\left(\Phi^{\otimes(k-1)} \otimes \mathbb{1} \otimes \Phi^{\otimes(N-k)}\right) \times \right. \\ \left.\left(\mathbb{1}^{\otimes(k-1)} \otimes \rho^{\theta_k} \otimes \mathbb{1}^{\otimes(N-k)}\right)\right) c_{\phi_k,\theta_k} = \\ \frac{1}{N}\frac{1}{L}\sum_{k=1}^{L}\sum_{\phi_k}\sum_{\theta_k} Tr_k\left(Tr_{\bar{k}}\left(\tilde{\Pi}_{\phi_k}^{k}\left(\Phi^{\otimes(k-1)} \otimes \mathbb{1} \otimes \Phi^{\otimes(N-k)}\right) \times \right.\right. \\ \left.\left.\left(\mathbb{1}^{\otimes(k-1)} \otimes \rho^{\theta_k} \otimes \mathbb{1}^{\otimes(N-k)}\right)\right)\right) c_{\phi_k,\theta_k} = \\ \frac{1}{N}\frac{1}{L}\sum_{k=1}^{L}\sum_{\phi_k}\sum_{\theta_k} Tr_k\left(\rho^{\theta_k}\, Tr_{\bar{k}}\left(\tilde{\Pi}_{\phi_k}^{k}\left(\Phi^{\otimes(k-1)} \otimes \mathbb{1} \otimes \Phi^{\otimes(N-k)}\right)\right)\right) c_{\phi_k,\theta_k} \end{aligned}$$

Since the partial trace is a positive trace preserving superoperator, the operator

$$\Pi_{\phi_k}^{k} = Tr_{\bar{k}}\left(\tilde{\Pi}_{\phi_k}^{k}.\left(\Phi^{\otimes(k-1)} \otimes \mathbb{1} \otimes \Phi^{\otimes(N-k)}\right)\right) \tag{8.28}$$

is a positive operator. Moreover, it is easy to verify that $\sum_{\phi} \Pi_{\phi}^{k} = \mathbb{1}$ so the operators $\{\Pi_{\phi}^{k}\}_{\phi}$ comprise a complete set of POVM elements acting on the k^{th} subsystem. Thus we have:

$$p_{forgery}^{average} = \frac{1}{L}\sum_{k=1}^{L}\frac{1}{N}\sum_{\phi_k}\sum_{\theta_k} Tr\left(\rho^{\theta_k}\Pi_{\phi_k}^{k}\right) c_{\phi_k,\theta_k}$$

and we have expressed the average probability of a cheating Bob causing a photodetection event with Charlie in terms of individual measurements on quantum signature states, without any assumption on the choice of the global measurement. This means that Bob can achieve the same cheating probability with collective measurements as with uncorrelated individual measurements. But then his optimal strategy is to use one best individual measurement for each state, which is the case we analysed earlier. This concludes our analysis which shows that any cheating probability achieved by Bob using a global measurement can be realized by independent single system measurements.

8.3.4 Active strategy – separable attacks

In this section we analyse Bob's forging probabilities in the case he employs an active, separable strategy. In active separable strategies, Bob is allowed to alter the states he sends to Charlie during the quantum signature distribution phase, but his malevolent activity is assumed to be equal for each quantum signature element state, and he also acts individually and identically on each element state. By altering the states he sends to Charlie, Bob can try to increase the probability to successfully forge a message later on. We will call the states Bob sends to Charlie the response states. Here, to guarantee security, we must take into account Charlie's multiport null-port photodetection events. For the k^{th} element of the quantum signature, which has a phase of θ_k, Bob has access to his copy of the quantum signature, along with the half pulse he received from Charlie. This can be represented by the state $|e^{i\theta_k}\sqrt{3/2}\alpha\rangle$ in total. In order to forge a message in the future, Bob will at some stage have to commit to an angle ϕ_k, which will comprise the forged private key. To select the best angle to commit to for the private key, Bob makes a generalized measurement on a fraction of the state $|e^{i\theta_k}\sqrt{3/2}\alpha\rangle$, and we allow this fraction to be anything between zero and unity [3]. We may assume that the measurement takes place before Bob sends a response state to Charlie within the multiport. This holds, without the loss of generality, because knowing the result of the measurement can only improve Bob's ability to select a response state that would increase his probability of successfully forging a message. The response state η^{θ_k,ϕ_k} may in general depend on both the actual phase value of the k^{th} quantum signature state and on Bob's measurement outcome. We note that in the case of a passive strategy, the response state will be $|e^{i\theta_k}\alpha/\sqrt{2}\rangle$. The forwarded, possibly altered response state is then interfered on Charlie's final multiport beamsplitter with Charlie's half of the k^{th} quantum signature state. One output arm (the multiport null-port) is measured for a photon count, and the output state of the other arm is stored by Charlie as the k^{th} quantum signature state. Please see Figure 8.1 for an illustration.

[3] Since the allowed measurement is a generalized measurement, it is actually superfluous to explicitly state that Bob is allowed to measure any fraction of the pulse. However it may serve the purpose of emphasizing we are not placing any restrictions on Bob's individual measurements.

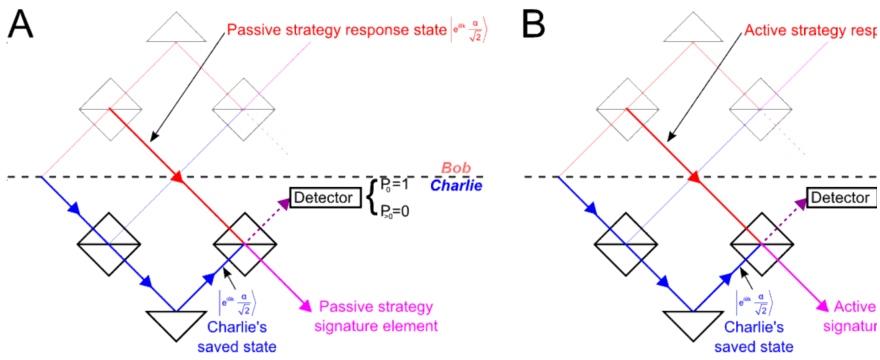

Figure 8.1. (A) In the passive strategy, Bob's response state is equal to Charlie's saved state. Consequently, the probability P of not detecting a photon at the multiport null-port is unity, whereas the probability of detecting one or more photons $P_{>0}$ is zero. (B) In an active strategy the probability P_0 is non-zero. It is, however, equal to the expected fidelity between the passive and active strategy signature elements.

The response state can be written in the most general P-representation form as $\eta^{\theta_k,\phi_k} = \int_{\mathbb{C}} P(\beta)|\beta\rangle\langle\beta|d^2\beta$ and the joint state of the Charlie's final beamsplitter is then:

$$\int_{\mathbb{C}} P(\beta) \underbrace{|\beta/\sqrt{2} - e^{i\theta_k}\alpha/2\rangle\langle\beta/\sqrt{2} - e^{i\theta_k}\alpha/2|}_{null-port} \otimes \underbrace{|\beta/\sqrt{2} + e^{i\theta_k}\alpha/2\rangle\langle\beta/\sqrt{2} + e^{i\theta_k}\alpha/2|}_{quantum\ signature\ element} d^2\beta \tag{8.29}$$

The probability of detecting no photons at the null-port arm is then:

$$Tr(|0\rangle\langle 0|. \int_{\mathbb{C}} P(\beta)|\beta/\sqrt{2} - e^{i\theta_k}\alpha/2\rangle\langle\beta/\sqrt{2} - e^{i\theta_k}\alpha/2|d^2\beta) = Tr(|e^{i\theta_k}\alpha/2\rangle\langle e^{i\theta_k}\alpha/2|.\eta') \tag{8.30}$$

where $\eta' = \int_{\mathbb{C}} P(\beta)|\beta/\sqrt{2}\rangle\langle\beta/\sqrt{2}|d^2\beta$. The state that Charlie will store as the quantum signature key is given with

$$\mathbf{D}(e^{i\theta_k}\alpha/2)\eta'\mathbf{D}^\dagger(e^{i\theta_k}\alpha/2), \tag{8.31}$$

where $\mathbf{D}(\cdot)$ denotes the displacement operator. The expression of the right hand side of equality (8.30) is sometimes referred to as the expected fidelity between the states $|e^{i\theta_k}\alpha/2\rangle\langle e^{i\theta_k}\alpha/2|$ and η'. Recall that, in the case of a passive strategy, the state η' would be exactly $|e^{i\theta_k}\alpha/2\rangle\langle e^{i\theta_k}\alpha/2|$. From this, it is easy to see that the probability of detecting a photon at Charlie's multiport null-port is equal to the expected fidelity between the quantum signature elements Charlie will store in the active and passive attack settings, respectively. This we have illustrated in figure 8.1.

The process of signature verification for each quantum signature element is a two-outcome measurement the outcomes of which correspond to the detector either registering a photon or not. If a bound of the trace distance between the average stored signature elements in the passive and active attacks can be guaranteed, then we can bound the difference of causing a photon detection event during signature verification for the active and the passive strategies. This is ensured by setting a rejection threshold on the multiport null-port photon event count.

Let r be the fraction of the quantum signature states, which have caused a photon event during signature distribution, where the quantum signature is of length L. Recall, we are assuming Bob is acting independently and identically for each key state so this statistic can be used to bound the value of $Tr(|e^{i\theta_k}\alpha/2\rangle\langle e^{i\theta_k}\alpha/2|.\eta')$ for an average key state. Let

$$x := 1 - Tr(|e^{i\theta_k}\alpha/2\rangle\langle e^{i\theta_k}\alpha/2|\eta'). \tag{8.32}$$

Then, by the Hoeffding's inequality we have that $P(|x - r| \geq \epsilon) \leq 2exp(-2\epsilon^2 L)$. Thus we have that

$$1 - Tr(|e^{i\theta_k}\alpha/2\rangle\langle e^{i\theta_k}\alpha/2|\eta') \leq r + \epsilon \tag{8.33}$$

except with probability $2exp(-2\epsilon^2 L)$. The expected fidelity has a well-known relationship with the trace distance:

$$D(|e^{i\theta_k}\alpha/2\rangle\langle e^{i\theta_k}\alpha/2|, \eta') \leq \sqrt{1 - Tr(|e^{i\theta_k}\alpha/2\rangle\langle e^{i\theta_k}\alpha/2|\eta')} = \sqrt{r + \epsilon}. \tag{8.34}$$

Thus, we have that the trace distance between the average response state and the state Bob would send in a passive strategy is less than $\sqrt{r + \epsilon}$ if the fraction of photon detection events was r except with probability $2exp(-2\epsilon^2 L)$. With probability $2exp(-2\epsilon^2 L)$ the trace distance between the response state and the passive strategy state may be above $\sqrt{r + \epsilon}$, but is always below unity. Thus, we can bound the average trace distance as follows [4]:

$$D(|e^{i\theta_k}\alpha/2\rangle\langle e^{i\theta_k}\alpha/2|\eta') \leq (1 - 2exp(-2\epsilon^2 L))\sqrt{r + \epsilon} + 2exp(-2\epsilon^2 L). \tag{8.35}$$

One can then ensure that the trace distance between the response state and the state corresponding to a passive strategy is arbitrarily small (except with the small probability $2exp(-2\epsilon^2 L)$), by selecting a rejection threshold r and a value ϵ. Then the trace distance approaches $\sqrt{r + \epsilon}$ exponentially quickly in the quantum signature length L. Let us denote the upper bound on the trace distance $\delta = (1 - 2exp(-2\epsilon^2 L))\sqrt{r + \epsilon} + 2exp(-2\epsilon^2 L)$.

Recall, in the passive attack the probability of not causing a photon detection event per quantum

[4] State verification, or even state tomography, of this type cannot ever guarantee that the verified state lies in some $\sqrt{r + \epsilon}-$ ball, with respect to the trace distance or any other reasonable measure, around the target state with unit probability. Thus, the effective trace distance above just averages between the individual average quantum state forging probability, given the measured state was within the $\sqrt{r + \epsilon}-$ball, and total protocol failure per average quantum element state, when the verified state was outside the $\sqrt{r + \epsilon}-$ball. Additionally, for reasonable values of ϵ, the probability of the response state being outside the $\sqrt{r + \epsilon}-$ball around the honest response state is negligible.

signature state was given with the following minimum cost:

$$p_{forgery} = \frac{1}{N} \sum_{\phi} \sum_{\theta} Tr(\Pi_{\phi} \rho^{\theta}) c_{\phi,\theta}.$$

For our cost matrix, Bob's optimal measurement was shown to be the square-root measurement, and this was dependant on the structure of the cost matrix c.

In the active attack, Bob has access to a larger amplitude coherent pulse then in the passive setting, as in principle he can measure Charlie's fraction of the signature states as well. Since we impose a restriction on Bob's activities by checking the multiport null-port count, in practice Bob will not be able to measure out all of the systems he receives, as he is forced to return a perhaps slightly modified variant of half of Alice's signature element, or Charlie's half of the coherent pulse, in order to pass Charlie's null-port test during signature distribution. To lower bound Bob's cheating probabilities we, however, assume that he can indeed use the entirety of the state he has received from both Alice and Charlie for the measurement. This is equivalent to giving Bob amplified versions of the quantum signatures. We will denote the probability of Bob not causing a photodetection event in a passive strategy with amplified pulses by

$$p_{forgery}^{amplified,\ passive} = \frac{1}{N} \sum_{\phi} \sum_{\theta} Tr(\Pi_{\phi} \rho_{amplified}^{\theta}) c_{\phi,\theta}, \tag{8.36}$$

where $\rho_{amplified}^{\theta} = |e^{i\theta}\sqrt{\frac{3}{2}}\alpha\rangle\langle e^{i\theta}\sqrt{\frac{3}{2}}\alpha|$.

The induced value $p_{forgery}^{active}$ is lower-bounded by the optimization of the minimum cost problem related to the one above, but where the entries of the cost matrix have been decreased by δ. The following derivation shows that the induced value $p_{forgery}^{active}$ deviates from $p_{forgery}^{amplified,\ passive}$ by no more than delta:

$$\begin{aligned}
1/N \sum_{theta} \sum_{phi} Tr(\Pi_{\phi} \rho^{theta})([C]_{\phi,\theta} - \delta) = \\
1/N \sum_{theta} \sum_{phi} Tr(\Pi_{\phi} \rho^{theta})([C]_{\phi,\theta}) - 1/N \sum_{theta} \sum_{phi} Tr(\Pi_{\phi} \rho^{theta})(\delta) = \\
1/N \sum_{theta} \sum_{phi} Tr(\Pi_{\phi} \rho^{theta})([C]_{\phi,\theta}) - 1/N \sum_{theta} Tr(\mathbb{1} \rho^{theta})(\delta) = \\
1/N \sum_{theta} \sum_{phi} Tr(\Pi_{\phi} \rho^{theta})([C]_{\phi,\theta}) - 1/N \sum_{theta} (\delta) = \\
1/N \sum_{theta} \sum_{phi} Tr(\Pi_{\phi} \rho^{theta})([C]_{\phi,\theta}) - \delta.
\end{aligned}$$

Thus we have the bound

$$p_{forgery}^{active} \geq p_{forgery}^{amplified,\ passive} - \delta. \tag{8.37}$$

For illustration purposes, the value $p_{forgery}^{amplified,\ passive} = 4.61 \times 10^{-3}$ for our experimental set-up, inducing a slightly reduced gap of approximately $g^{amplified} = 7.13 \times 10^{-4}$, which is slightly worse than the passive strategy value g of approximately 8×10^{-4}. The calculation of $g^{amplified}$ is performed analogously to the calculation of g given in previous sections, with the same cost matrix, but amplified measured states in the minimum cost problem 8.16.

The overall probability of Bob forging is given again by the Hoeffding's inequalities as

$$P_{Bob\ Cheats} \leq 2\exp(-\frac{2}{9}(g^{amplified} - \delta)^2 L),$$

which again approaches zero exponentially quickly in the signature length L, as long as we ensure that $\delta < g^{amplified}$. Just as an illustration we performed the calculations estimating the key lengths required to ensure δ is below the amplified gap $g^{amplified}$, and the required key lengths are at the order of 10^9. The efficiency of this protocol, with the used parameters could benefit from an improvement, and this we discuss in Section 9.1.

In this analysis for active forging attacks, we assumed that the detectors at the multiport null-port were perfect, whereas imperfections would increase the required key length to ensure the estimate of δ. However, we have also used lower bounding results which are not tight. Here, we briefly consider the effects of imperfect devices on the security parameters. First, taking into account the known detector losses one can still work out the required rejection threshold r and the required signature length L to ensure $g^{amplified} > \delta$. The losses will make the required threshold lower, and the accompanying signature length L longer, when compared to the values obtained in the ideal case previously. Nonetheless, arbitrarily small values of δ can still be obtained efficiently in the signature length. Additionally, the differential losses occurring within the multiport would also cause Bob and Charlie to have differential sensitivities to cheating. In our experiment, the measurements of the cost matrix C given in sections 8.3.2 and 7.2, from which the gap g is calculated were performed on the party with the lower overall losses. The party with the losses lowered by a multiplicative constant c would realize a guaranteed value of the gap $g' = cg$.

In our case $c > 1/4$ (see Figure 7.5). Finally, we briefly address the question of the protocol's robustness with respect to the rejection threshold r, which limits the acceptable multiport null-port photon counts. In the presence of dark counts this threshold may be breached, even when all parties are honest. In our system the raw dark count probability per emitted pulse per detector stands at approximately $p(dark) = 3.2 \times 10^{-6}$. To take this into account, the baseline threshold r, chosen to achieve the required levels of security against active strategy forging, should be increased by the value of $p(dark)$. The realized security levels are not jeopardized as no cheating response state Bob may send can reduce the dark count rate [5]. The dark count rate is limited by the detector and can only possibly increase from the baseline level, which is realized in as passive strategy. Analogous arguments should hold if additional causes for a photon detection event not due to a cheating response state, such as interferometric visibility and background count rate, are considered. The honest setting rejection probability for such a setting can again be shown to

[5] This holds barring the recently demonstrated attacks by the group of Vadym Makarov where they actually do reduce the dark counts by effectively burning the photodetector using strong light, as presented at QCMC 2012.

vanish exponentially quickly in terms of the signature length L as

$$\epsilon_{robustness}^{multiport} \leq \exp(-2r^2 L). \tag{8.38}$$

8.3.5 Active strategy – coherent attacks

Here, we give a plausibility argument that coherent, or any type of general strategy Bob may employ, does not improve Bob's forging probabilities when compared to the separable cheating strategy, discussed in the previous section. The technique we suggest to show this is analogous, albeit more involved, to the one used in the proof of security against repudiation for the coherent attacks. We shall consider two types of fictitious protocols – games – which are obtained by modifying the original protocol, in the context of the two types of strategies by Bob: individual and coherent. We will denote the original protocol by **O**. A modified original protocol, which we call a sequential, delayed, with disclosure protocol, we will denote **SDwD**. Finally, we will use a modified protocol called a delayed protocol, denoted D. In the **SDwD** protocol, the states sent by Alice are accumulated halfway within the multiport: Bob receives all the original quantum signature element states, along with the "half" of the pulse from Charlie, and Charlie accumulates all his "half" pulses. This constitutes the "delay" in the designation of this modified protocol. From here the protocol continues sequentially: at the k^{th} step, Bob sends the k^{th} response state, Charlie interferes it with his corresponding "half" pulse and obtains the corresponding null-port measurement outcome. At this point, Bob chooses, or commits to, a private key element (the phase angle) based upon the measurement of whatever system he may have. Without the loss of generality this commitment/measurement could have taken place at the instance of Bob generating the response state. Bob then declares his guess of the private key element, and Charlie proceeds to perform the verification for this signature element. This constitutes the "sequential" attribute in the protocol designation. Finally, Alice, who is an honest player in this modified game, at this point reveals the actual angle which she encoded in the k^{th} pulse to the cheater Bob. Note, that this happens after the verification for this pulse has been carried out. This procedure is sequentially repeated for all signature elements, and Bob wins the game, an event we shall call $Bob\ cheats$, if he managed to pass both the null-port and the verification thresholds. The modified protocol D just introduces the change of first accumulating the states within the multiport (the "delay" explained above), before continuing, and is otherwise identical to the original protocol "O". Concerning Bob's activities, we distinguish a separable strategy corresponding to individual identical activities, denoted S, and a coherent strategy C. In a separable strategy S, Bob chooses a response state, and commits to a to-be-declared private key element ("best guess" phase) by measuring the quantum signature states individually and an identical strategy is applied for each signature element. Also, the response states are not entangled with each other. The probability $P(Bob\ cheats|O,S)$, *i.e.* the probability of Bob successfully forging using a separable attack in the original protocol, is the value computed in the previous section. In a coherent strategy C, Bob is not restricted in any way, aside from the protocol specifications, which would otherwise cause an implicit protocol abort. An example of this would be Bob's failure to choose a phase angle and declare it to Charlie at any step of the SDwD protocol. Our goal is to prove the following

sequence of (in)equalities:

$$P(Bob\,cheats|O,S) \leq P(Bob\,cheats|O,C) \quad \text{(a)}$$
$$P(Bob\,cheats|O,S) = P(Bob\,cheats|SDwD,S), \quad \text{(b)}$$
$$P(Bob\,cheats|SDwD,S) = P(Bob\,cheats|SDwD,C), \quad \text{(c)}$$
$$P(Bob\,cheats|SDwD,C) \geq P(Bob\,cheats|D,C), \quad \text{(d)}$$

and finally

$$P(Bobcheats|D,C) \geq P(Bobcheats|O,C). \quad \text{(e)}$$

The sequence of inequalities (b) - (e) then shows

$$P(Bob\,cheats|O,S) \geq P(Bob\,cheats|O,C), \tag{8.39}$$

which combined with the inequality (a) yields the desired claim,

$$P(Bobcheats|O,S) = P(Bobcheats|O,C). \tag{8.40}$$

The first claim (a) is trivial, as a separable strategy is a special case of coherent strategies. The claim (b) is relatively easy as well: due to the separable nature of Bob's attack, neither having all the states at his disposal simultaneously, nor the "sequential" modification play a role. Since the disclosure of the outcomes and the angles comes after Bob's activity per element state, and since the phases in the quantum signature states are independently and uniformly chosen at random, this information cannot help Bob either. To first sort out the obvious claims, we note that the claim (e) is trivial as well. Bob can, in the delayed setting, run every strategy he can run in the original setting, thus delay can only help. This confirms (e). For the remainder of this section we will thus be focusing on claims (c) and (d). We start with the claim (c),

$$P(Bob\,cheats|SDwD,S) = P(Bob\,cheats|SDwD,C), \tag{8.41}$$

for which we give a plausible motivation, but strictly speaking not a proof. In the (SDwD, S) setting, Bob accumulates the states he receives within the multiport (from both Alice and Charlie) and then individually acts identically, using a most general physical procedure allowable by quantum mechanics, on each individual state, sequentially producing a response state along with private key element. At each step, this is followed by two measurements by Charlie (null-port and the verification-constituting measurement) and a total disclosure of the originally encoded angle and the measurement outcomes. In contrast, in the (SDwD, C) setting, Bob is allowed to perform any global operation on the states he has, and any ancillary system he may wish to use, but again he has to, at each step, choose a response state and an angle (the private key element). As every response state is processed in run-time, Bob has nothing to gain by using response states which

are entangled with his remaining subsystem. Note that, since he gets the disclosure information, whatever system he would have following Charlie's measurements, Bob can simply reproduce post disclosure at each step. We now focus on the first response state and angle Bob will generate, and see whether Bob can increase his probabilities of favourable outcomes of Charlie's measurements by using global operations. Note, that the measurement Charlie performs for the first state depends only on the angle encoded within the first quantum signature element sent by Alice. This is independent of all subsequent signature states. Thus, no response state, generated based on information Bob may gain by globally measuring the entirety of his system, which includes the entire amplified quantum signature, can influence the outcomes of Charlie's measurements for the first system in a way that benefits Bob, when compared to the separable strategy. Next, we check whether a coherent strategy starting at the first step can augment his probabilities at later stages. As noted, since all measurement outcomes are disclosed after Charlie's measurement at each step, whatever state Bob remains with using in the coherent strategy after the measurement, Bob can generate using the disclosed information in a separable strategy. Hence, at each step, his strategy for that particular step may as well be a separable one, as the full disclosure at each step nullifies any advantage he might have gained by using a coherent strategy. Coherent strategies *seem* to give Bob no advantage in this setting and claim (c) seems *plausible*. To iterate, this however, does not constitute a proof, but maybe a proof idea. To finalize our argument we analyse claim (d),

$$P(Bob\,cheats|SDwD, C) \geq P(Bob\,cheats|D, C). \tag{8.42}$$

The only advantage the protocol variant D may hold for Bob is that he need not commit to a particular private key element angle in run-time, but rather can do a global measurement on his system later. Note that if we assume that there exists no communication between Bob and other parties until Bob wins or loses the game (*i.e.* cheats or get caught) in the D variant of the protocol then delaying the measurement cannot increase his cheating probability. In the D protocol Bob does obtain the information whether the multiport null-port threshold has been breached. However, since both the verification and the multiport null-port thresholds violations constitute Bob losing, his measurement strategy should not be conditional on whether he passed the multiport threshold. Provided this holds, in the D protocol we may assume that Bob measures his system at any point up to the moment when he sends away his last response state. Any measurement Bob may perform in the D protocol can be realized by a large unitary acting on the entirety of his system and a sufficient amount of ancillary systems, followed by single system measurements. This can be seen as a consequence of the Naimark dilation theorem. The single system measurement outcomes will give Bob's choices of private key elements. However, Bob, if he employs a coherent strategy, may perform the same map in the SDwD setting as well, and measure the angle-carrying subsystems sequentially as he is required by the protocol. Bob can thus obtain the same cheating success probability in the SDwD setting as in the D setting by simply ignoring the information he gets from the disclosure in the SDwD setting. But then clearly, gaining additional information can only help Bob, and we have our inequality (d). More formal variants of the proof of claims (c) and (d) we leave for further research. This concludes

our argument.

8.4 Technical results

In this section we prove the technical lemmas and present mathematical statements we used in the security of analysis of the experimental system.

8.4.1 Hoeffdings inequalities

Here, we briefly state Hoeffding's inequalities and explain how they are used in the security analysis. We are presenting a restricted version of these inequalities, which is more directly applicable to our setting.

Lemma 26. *Let $X_1, \ldots, X_L$ be independent random variables each attaining values 0 or 1. Let $\bar{X} = 1/L \sum X_i$ be the empirical mean of the variables, and let $E(\bar{X})$ be the expectancy of the empirical mean. Then we have:*

$$P(\bar{X} - E(\bar{X}) \geq t) \leq \exp(-2t^2 L) \tag{8.43}$$

$$P(|\bar{X} - E(\bar{X})| \geq t) \leq 2\exp(-2t^2 L). \tag{8.44}$$

In the case of the analysis of Bob's forgery probabilities, we compute the minimal probability of obtaining a photon detection event on the multiport p_{cheat}, which defines a sequence of L random variables as in the statement of the theorem above. Then we set a threshold at $s_v M$, and ask ourselves what the probability of an empirical mean of the random variables given above diverging from the expectancy by more than $p_{cheat} - s_v$ is. Note that this is the requirement for the forgery to be accepted. This corresponds to the second inequality, as the empirical value needs to be below the mean/expectancy, and we are finished as soon as we introduce the values corresponding to our scenario, as was done in the presented analysis.

Robustness is calculated similarly, with the random variables defined by $p_{original}$ and we ask what the probability of getting an empirical mean which violates the threshold s_a is. This problem corresponds to the first of the two listed Hoeffding's inequalities. However, we need to take into account that both Bob and Charlie could reject in the honest setting, so to upper bound this value, using the union bound, we add the two obtained probabilities.

8.4.2 Trace distance and effects

This section corresponds to the section about repudiation probability for realistic systems. For the trace distance between two states $D(\sigma, \rho)$ we have the following property:

$$D(\sigma, \rho) \geq \frac{1}{2} \sum_x |Tr(\Pi_x \rho) - Tr(\Pi_x \sigma)| \tag{8.45}$$

for any set of POVM elements $\{\Pi_x\}_x$. A characterisation of the trace distance is given by taking the maximum of the right-hand side of the equation above over all POVMs. In our case ρ is the perfectly symmetric system with respect to Charlie's and Bob's subsystems, and σ whatever we actually produce in the lab, and the POVM is the four outcome POVM giving the possible outcomes of photon detection on their individual subsystems. Assume that malevolent Alice's target output is that Bob accepts and Charlie rejects, hence she wishes to maximize the probability of the outcome $(0,1)$. Let $p_x := Tr(Pi_x\rho)$, and $q_x := Tr(Pi_x\sigma)$ for $x \in \{(0,0),\ldots,(1,1)\}$. So we have

$$D(\sigma,\rho) \geq \frac{1}{2}\sum_x |p_x - q_x| = \frac{1}{2}\left|p_{(0,1)} - q_{(0,1)}\right| + \frac{1}{2}\sum_{all\ but\ (0,1)} |p_x - q_x| \tag{8.46}$$

Note that if $\left|p_{(0,1)} - q_{(0,1)}\right| = \epsilon$ therefore $\sum_{all\ except\ (0,1)} |p_x - q_x| \geq \epsilon$, so we have:

$$D(\sigma,\rho) \geq \left|p_{(0,1)} - q_{(0,1)}\right|. \tag{8.47}$$

From this we have the claims given in Section 8.2.3.

8.4.3 Minimum cost measurement problem for special cost matrices

Here, we prove the technical lemmas from the previous sections of this chapter. First we standardize the notation.

- Received states: These are the states Bob receives from Alice and Charlie jointly, so if the amplitude of the individual states of the unperturbed keys is α, Bob will receive the states $|v_k\rangle = |\exp(2k\pi I/N)\sqrt{3/2}\alpha\rangle$.
- Standard basis: For the states $|v_k\rangle := |e^{2k\pi I/N}\frac{\alpha}{\sqrt{3/2}}\rangle$ one can show that the states

$$|b_k\rangle := 1/\sqrt{\lambda_k N}\sum_{l=0}^{N-1} exp(-2kl\pi I/N)|v_l\rangle \tag{8.48}$$

  form an orthonormal basis, where λ_k are the eigenvalues of the Gram matrix of these states, and λ_k also comprise the diagonal, belonging to the diagonal matrix $\sum_{k=0}^{N-1}|v_k\rangle\langle v_k|$. The symbol I here denotes the imaginary unit. In this basis the states $|v_k\rangle$ have the following expansion:

$$|v_k\rangle := 1/\sqrt{N}\sum_{l=0}^{N-1} exp(2kl\pi I/N)\sqrt{\lambda_l}|b_l\rangle \tag{8.49}$$

  and, in particular, the vector $|v_0\rangle$ has in this basis all entries positive.

- The unitary characterising the symmetry of the system: U such that $U|v_k\rangle = |v_{k+1}\rangle$, index addition taken modulo N, indexing starting with zero. In the standard basis this unitary is

diagonal: $U = \sum_{l=0}^{N-1} \exp(2\pi li/N)|b_k\rangle\langle b_k|$.

- With DFT we denote the discrete Fourier transform matrix of (implicit) size N , defined element-wise by $[DFT]_{p,q} = \exp(-2\pi Ipq/N)$ for $p = 0...N1, q = 0, ..., N1$.

The properties we state above without proof are a direct consequence of the symmetricity of the sets of states we consider. For a more general theory of the properties of symmetric sets of states and their relationship with the discrete Fourier transform from which all the results follow, we refer the reader to the results we present in Chapter 10 of this thesis.

Claim (Helstrom criteria for symmetric states and positive circulant cost matrices) *If the input states are symmetric, and the cost matrix is circulant, then the Helstrom condition 3 holds for the square-root measurement.*

Proof sketch:
We have the risk operators defined as:

$$W_i = \sum_j c_{i,j}|v_j\rangle\langle v_j| = \sum_j c_{i,j}U^j|v_0\rangle\langle v_0|U^{-j}. \tag{8.50}$$

If the cost matrix $C = [c_{i,j}]_{i,j}$ is circulant we have that:

$$U^k W_0 U^{-k} = W_k. \tag{8.51}$$

The Lagrangian operator is defined as

$$\Gamma = \sum_i \Pi_i W_i. \tag{8.52}$$

The square-root measurement is defined by the operators

$$\Pi_i = \Phi^{-1/2}|v_i\rangle\langle v_i|\Phi^{-1/2}, \tag{8.53}$$

where $\Phi = \sum_i |v_i\rangle\langle v_i| = \sum_i U^i|v_0\rangle\langle v_0|U^{-i}$.

We will use the following property:
Lemma 27. *For any square matrix A we have that:*

$$\sum_i U^i A U^{-i} = NA' \tag{8.54}$$

where A' is the diagonal matrix containing the main diagonal of A.

Proof:
Let $|\omega_l\rangle = \sum_k \exp(2\pi ikl/N)|b_k\rangle$. The ket $|\omega_l\rangle$ in the standard basis contains the main diagonal of the matrix U^l. Then it is easy to see that for all square matrices A we have that $U^l A U^{-l} = A \circ |\omega_l\rangle\langle\omega_l|$, where $\circ$ denotes the Hadamard (Schur, point-wise) matrix product, which is distributive

with respect to matrix addition. Then we have

$$\sum_i U^i A U^{-i} = \sum_i A \circ |\omega_i\rangle\langle\omega_i| = A \circ \sum_i |\omega_i\rangle\langle\omega_i| \tag{8.55}$$

Using the properties of the sums of roots of unity we have that $\sum_i |\omega_i\rangle\langle\omega_i| = M\mathbb{1}$ where $\mathbb{1}$ is the identity matrix. Hence, we have our claim, as Hadamard-multiplying any matrix with the identity simply eliminates all off-diagonal elements.□

So we have that $\Phi = \sum_i U^i |v_0\rangle\langle v_0| U^{-i} = N|v_0\rangle\langle v_0| \circ \mathbb{1}$. Thus Φ is diagonal, and by the form of the ket $|v_0\rangle$ it simply collects the eigenvalues of the Gram matrix of the input states across the diagonal. But then $\Phi^{-1/2}$ contains the inverses of the roots of the eigenvalues λ_k across the diagonal. Since it is also diagonal U and Φ and $\Phi^{-1/2}$ commute, so we have that: $\Pi_k = U^k \Pi_0 U^{-k}$ and for the Lagrangian we have that:

$$U^k \Gamma U^{-k} = \Gamma. \tag{8.56}$$

We will also use a slightly more involved lemma, which generalizes Lemma 10.7:

Lemma 28. *For any square matrix A, and a sequence of N complex numbers $\{c_i\}_{i=0}^{N-1}$ we have that:*

$$\sum_i c_i U^i A U^{-i} = A \circ B \tag{8.57}$$

where B is a circulant matrix, and its first row is given with $DFT.\,[c_0, \ldots, c_{N-1}]^T$, i.e. *the discrete Fourier transform of the vector with entries c_i.*

Proof sketch:

Similar to the proof of the simpler lemma, with realization that $\sum_i c_i |\omega_i\rangle\langle\omega_i|$ a circulant matrix, and its first row is given with $DFT.\,[c_0, \ldots, c_{N-1}]^T$.□

This lemma is applied on the risk operator W_0 to obtain the following:

$$W_0 = |v_0\rangle\langle v_0| \circ B, \tag{8.58}$$

where B is a circulant matrix the first row of which comprises the eigenvalues of the cost matrix. To see this simply note that the cost matrix is circulant, and the eigenvalues of a circulant matrix are given by the DFT of the first row of the matrix. We need to show that $\underbrace{(W_i - \Gamma)\Pi_i = 0}_{eq.1} = \underbrace{0 = \Pi_i(W_i - \Gamma)}_{eq.2}$. Because of the symmetries we have that $(W_i - \Gamma)\Pi_i = 0$ if and only if $(W_0 - \Gamma)\Pi_0 = 0$ and the analogous holds for the second equality above.

To prove equality $eq.1$ we first show that the following holds:

$$W_0 \Pi_0 = \Gamma \Pi_0. \tag{8.59}$$

Using all the properties listed above, this is not a complicated task by explicitly inserting in all the operators and expanding the expressions. We omit this exercise. Analogously, one shows equality $eq.2$

$$\Pi_0 W_0 = \Pi_0 \Gamma. \tag{8.60}$$

holds, and we have our claim. □

Lemma 29. *If the cost matrix if positive, symmetric and circulant then the first and second Helstrom criteria are satisfied for the minimum error (square-root) measurement for our problem.*

Proof sketch:

Since W_i is a sum of positive operators with positive weights it is a positive operator, and specially Hermitian. The operators Π_i are positive as they are POVM elements, hence positive and Hermitian as well. Thus, we have

$$\Gamma = (\Gamma^\dagger)^\dagger = ((\sum_i \Pi_i W_i)^\dagger)^\dagger = (\sum_i W_i^\dagger \Pi_i^\dagger)^\dagger = (\sum_i W_i \Pi_i)\dagger \tag{8.61}$$

and the first and second Helstrom conditions are equivalent. Therefore, it suffices to show that:

$$\sum_i W_i \Pi_i = \sum_i \Pi_i W_i. \tag{8.62}$$

We have that

$$\sum_i W_i \Pi_i = \sum_i U^i W_0 \Pi_0 U^{-i} = N(W_0 \Pi_0) \circ \mathbb{1}$$
$$\sum_i \Pi_i W_i = \sum_i U^i \Pi_0 W_0 U^{-i} = N(\Pi_0 W_0) \circ \mathbb{1}$$

Hence, the lemma we are proving holds if and only if $(W_0\Pi_0)$ and $(\Pi_0 W_0)$ have equal diagonal elements. We have shown that

$$W_0 = |v_0\rangle\langle v_0| \circ B \tag{8.63}$$

where B is a circulant matrix where the first row comprises the eigenvalues of the cost matrix. Note that for the square root measurement we have the following property:

$$\Phi^{-1/2}|v_0\rangle\langle v_0|\Phi^{-1/2}|b_k\rangle = 1/\sqrt{N}\sum_l |b_l\rangle \tag{8.64}$$

Thus we have for the k^{th} diagonal element of $(W_0\Pi_0)$:

$$1/\sqrt{N}\langle b_k|(|v_0\rangle\langle v_0| \circ B)(\sum_l |b_l\rangle)$$

which is the sum of the elements of the representation of the bra $\langle b_k|(|v_0\rangle\langle v_0|\circ B)$ in the standard basis scaled by $1/\sqrt{N}$. For the k^{th} diagonal element of $(\Pi_0 W_0)$ we get

$$1/\sqrt{N}\sum_l \langle b_l|(|v_0\rangle\langle v_0|\circ B)|b_k\rangle$$

Since the matrix $|v_0\rangle\langle v_0|$ is real in the standard basis, these two expressions are equal if B is real and symmetric. Recall, the matrix B is the circulant matrix comprising the eigenvalues of the cost matrix. These are real if and only if the cost matrix is symmetric, which holds by the assumption of the lemma. The symmetricity of B is a consequence of the cost matrix comprising real elements. Thus the lemma holds. □

Chapter 9

Discussion: future of QDS

Here we put the QDS protocol in the context of modern quantum and classical cryptography, and attempt to assess how this protocol may evolve further. New schemes for QDS are proposed.

Quantum digital signatures were presented in 2001, and the global interest in this functionality has been modest, compared for example, to the success of QKD. Given the ever-growing importance of digital signatures in modern global e-business, it is curious why this idea until now did not assume a more prominent role on the quantum information processing stage. In this section we address, analyse, and attempt to resolve some of the potential issues with quantum digital signatures. We begin by inspecting the details of our experimental implementation of the original proposal of QDS, pinpointing potential impediments on the this "microscopic level" and proposing how the issues may potentially be resolved. Following this, we focus our attention more globally on the issues arising in QDS *by proposal design*, meaning issues which cannot be resolved by more elaborate ways of actually implementing the theoretical proposals. Finally, we take one final step back, and compare the performance and requirements of QDS relative to other successful cryptographic primitives.

9.1 Quantum digital signatures as on the experimental table

Resilience against most general attacks The analysis of security of the presented quantum signatures scheme did not include formal proofs of security of most general types of attacks. The security analysis of the most general types of forging attacks only gives a plausability argument that the system should be secure. While at this point we do not offer a resolution of this issue, we do believe that this can be resolved. With faith in this, we continue on with the analysis of more subtle and general issues with QDS.

Cost per signed bit Assuming the presence of an abundance of reliable quantum memory and maintained stability of our system, perhaps the largest issue, which is not satisfactory, is the efficiency of our realization. The experimental system we have built is based on the technology developed for high bit-rate QKD. In QKD, the improvement of the raw throughput of coherent pulses leads to a higher final bit-rate (assuming no increase in noise), thus more is necessarily

better. In the case of QDS, every pulse is assumed to be stored in quantum memory. The efficiency we discuss here is the following: how much of the most expensive resource do we need to spend per signed bit given a fixed desired level of security. In principle the two resources we may consider are (quantum) memory and time. In the presented protocol, these two resources are the same, unlike in the QKD scheme, where the final stored key will in practice be much shorter than the number of quantum states emitted.

But, as we shall show in the next section, the QDS protocol could potentially be tweaked so that the memory requirements are significantly reduced in quantity, and in particular, the need for quantum memory could be circumvented completely. Thus, we will simply be concerned with the number of pulses Alice needs to send to reach the desired levels of security. As we have shown in the first in-depth quantitative security analysis of a quantum digital signatures implementation presented in this work, the central figure characterising the security of the system, was the value g. This value appeared in three different incarnations in the cheating success formulae for both passive and active attacks. This value was derived directly from experimental data, and the relevant formulae were:

$$\begin{aligned} \epsilon_{forging} &= P_{Bob\ cheats} \leq 2\exp(-\frac{2}{9}g^2 L). \\ \epsilon_{forging}^{active} &\leq 2\exp(-\frac{2}{9}(g^{amplified} - \delta)^2 L) \\ \epsilon_{robustness} &\leq \exp(-\frac{2}{9}g^2 L) + \exp(-\frac{4}{9}g^2 L) \end{aligned} \tag{9.1}$$

and they upper bounded the forging probabilities in the passive and active strategies, and the robustness of the system in terms of the lengths of the quantum signatures used [1].

The value g was defined as $g = p_{forgery} - p_{original}$, where the values $p_{forgery}$ and $p_{original}$ were computed either directly from the cost matrix C, or by finding the minimum cost of a minimum cost measurement problem of distinguishing coherent states comprising the quantum signatures, with the cost matrix C. The value $g^{amplified}$ was defined in a similar way and it also depends on the same cost matrix C. For details see Section 8.3.1, and 8.3.4.

The values of g we present, while they do allow the generation of signatures of arbitrary levels of security exponentially quickly, the overhead for reaching interesting values of the cheating probabilities is formidable. In particular, let us consider formula 9.1 and the computed value $g \approx 10^{-4}$. Note that g appears squared in this formula, so the order of g^2 we are dealing with is at the order of 10^{-8}. From this it is obvious that, to reach non-trivial bounds of the cheating probability (*i.e.* below unity) the quantum signature length has to be of at least the order 10^8. These are long keys indeed, and even considering raising the repetition frequency of the electrical equipment driving the laser pulses to the gigahertz regime (from the current values of cca. 100 MHz) would render the minimal time required to emit the signatures for a single bit to be at the order of a second. In practice this is unsatisfactory, and the problem is not likely to be solved by pushing the technological limits of the repetition rate. The resolution here is magnifying the

[1] As we have explained in Chapter 8, the repudiation probability is below these values, but also depends on g.

value of g, and consequently $g^{amplified}$.

As it is defined, g depends on three parameters: the number of phases used N, the amplitude α and the experimentally determined cost matrix C (for details see sections 8.3.2 and 8.3.1). The cost matrix C itself depends on N, the number of phase encodings used, α and the actual experimental implementation of the system. We can attempt to increase g by theoretical and experimental means. From the theory point of view, even for a fixed and idealized experimental set-up, it is not understood for which values of N and α is g maximized. This problem could be addressed by employing numerical techniques if noise and losses are included in the model of the system, and for the ideal case even analytic results may be possible.

The numerical testing we have performed suggests that we may not have been working with the ideal values of N and α [2], and a more careful analysis followed by experimental confirmation may render improved values of g, albeit, the improvements of g we see in simulations do not reach a whole order of magnitude. To further increase the value of g we must inspect the aspects of the experimental set up which influence it.

More than likely, the main culprit for the low values of g are the effective low amplitudes of the states which get compared and Bob's (Charlie's) interferometers, 7.1. By effective amplitudes we mean the amplitudes that eventually enter the verification and authentication processes, after approximately 13 dB of losses. The effective amplitudes directly influence the overall magnitude of the values in the matrix C, and more importantly the relative differences which occur between the entries of C which correspond to an ever increasing discrepancy between the declared and the real angle. In particular, raising these effective values of α, relative to a fixed initial amplitude emitted by Alice could significantly improve the magnitude of g. There are two immediate reasons contributing to the effective reduction in the amplitudes, resulting from experimental imperfections, and *by construction* of the experiment.

Amplitude loss due to experimental imperfections These come from the losses predominantly in the multiport, estimated at 13 dB. As we have mentioned, a part of the loss is due to imperfect optical fibre splicing, air-gaps and the losses specific to other optical components we have used. Improvement at any of these would increase the effective amplitudes, and augment the value of g. The air-gaps could be replaced by comparatively lower loss fibre stretchers. However, fibre stretchers work as the name suggests by mechanically stretching the optical fibre. As the amount by which a given length of optical fibre may be stretched is low, long lengths are required to achieve the same levels of path length control as may be obtained with an air-gap. This increased length of the fibre could lead to a greatly reduced overall stability of the system, which may require greater periods of time spent on re-tuning the system, but could increase the value of g.

[2] The preliminary simulations we performed did not consider a sufficiently realistic noise model to make a firmer statement.

Amplitude loss by construction If we focus our attention at how the state comparison is performed at, say Bob's verification arm 7.1 we see that the state comparison is performed by modulating the reference pulse, which is then interfered with the signal pulse on the beamsplitter cube. But, the signal and reference pulses are both split on a simple beamsplitter cube, and the interference between the signal pulse and the modulated reference pulse is only achieved on one half of the coherence lengths of the two pulses, and the other two halves do not interfere. The non-interfering photon counts are gated away in post processing. This means that approximately $1/2$ of the amplitude is lost, by construction.

While the amplitude loss due to experimental imperfections cannot be reduced to arbitrary levels, the $1/2$ loss due to construction can certainly be improved upon. By changing the verification/authentication method and by improving on the initial implementation of the system, more favourable values of g may be achieved. As noted, even modest improvements in g may lead to substantially lower overheads, as g^2 appears in the actual expressions quantifying the security levels.

9.2 Quantum digital signatures on paper

Regardless of the method used to realize the QDS based on the idea of quantizing Lamport's scheme, there will always be two vital components which are very challenging to address in practice or experimentally. These are the quantum states comparison[3], and quantum memory. The central idea to the proposed schemes to QDS is to use quantum systems and the properties of quantum mechanics to ensure the "no forging" security requirement. Thus, the quantum signatures are quantum states, and they have to be distributed securely and kept in quantum memory until needed. Both of the latter two requirements are stepping stones and in this section we will address the latter one, leaving the issue of distribution for the next section of this chapter. Quantum state comparison is vital to this scheme to ensure security against repudiation, a security property no less essential than no forging. Implementing quantum state comparison is very difficult, especially given that the parties, which need to compare the states have to be assumed to be separated by real-world large distances. In this section we address the issues of quantum memory and state comparison and propose possible resolutions.

9.2.1 Quantum memory

While realizing the multiport which works over large distances may be very difficult to do with current technology, the memory requirement we need is even further away from reality. The quantum memory we need which can store billions upon billions of quantum modes making up the quantum signatures, and keep them without significant noise or loss over years, days, or even seconds with high reliability, is way beyond our technical abilities, certainly for years to come. Arguably, getting rid of the quantum memory requirement would be the biggest improvement

[3]More accurately, what we require is a quantum states "symmetrization", achieved with the multiport in our scheme, or a type of a SWAP test as in the original QDS proposal [2].

we could make to this protocol. The modification of the protocol which potentially achieves just this was invented and originally proposed by Dr. Erika Andersson, and the modification we need to introduce is simple. Instead of storing the quantum signatures, Bob and Charlie *just measure them*.

We present the simplest variant of this proposal in Protocol 15 and offer a proof sketch of security for the ideal case.

Recall, unambiguous state discrimination (USD) of a set of N states $S = \{\rho_i\}$ resolves the following problem. Given an unknown state from the finite set S, which may occur with some probability p_i perform a measurement which has the following $N + 1$ possible outcomes: $\{i\}_i$, denoting which state was measured, and additionally "?" denoting the procedure could not determine which state was measured. The process, however, has the following property: if an unambiguous outcome (in $\{i\}_i$) is obtained, then the measurement is guaranteed to have determined the state correctly. Thus, unlike minimum-error measurements we used in the security analysis of the experimentally realized QDS protocol, an USD measurement is not allowed to make a mistake, but it is allowed to produce an outcome "I don't know". The success probability of a USD procedure for a fixed set of states and occurrence probabilities is loosely defined as one minus the probability of obtaining the ambiguous outcome "?" [4]. USD procedures can succeed (*i.e.* not fail with a unit probability) if and only if the set of states S contains pure and linearly independent states. For more on this topic we refer the reader to Chapter 10, Section 10.1.2.1.

Now, we sketch out the proof of security for this protocol, under the same assumptions as in the case of the original protocol: that the quantum channels to the multiport are not accessible to potential forgers, and that the classical communication between Bob and Charlie is secure from Alice's tampering. As in the protocol with quantum memory, attacks on the channels leading to the multiport are usually outside the scope of theory of public-key cryptography, and in our case they would demolish security, by a key-swap attack. How and if this assumption may be dropped by employing a type of quantum message authentication we will discuss presently.

We now sketch the argument for the two properties of interest: no forging and no repudiation.

9.2.1.1 Security against forging

In order for, say Bob to forge a bit with Charlie, he has to reproduce the string Charlie generated in the distribution phase by measuring the elements of the quantum signatures, for the positions where Charlie obtained an unambiguous outcome.

[4] In the most general case, the probability of this outcome may depend on the state measured, which itself may have occurred with a non-uniform probability. In this case one may consider the average success probability.

Protocol 15 Quantum digital signatures with no quantum memory

1. To sign a single bit (message $m = 0$ or 1) in the future, Alice generates two sequences $PrivKey_0 = (b_1^0, \ldots b_L^0)$ and $PrivKey_1 = (s_1^1, \ldots s_L^1)$ of L randomly chosen angles from the set of signs $\{-1, 1\}$, so

$$b_k^m \in \{-1, 1\}. \tag{9.2}$$

   The pair $(m, PrivKey_m)$ is called a private key pair for message m.

2. Alice then generates two copies of a sequence of coherent states $QuantSig_0 = (\rho_1^0, \ldots, \rho_L^0)$ with the coherent phase of the sign matching the signs in the sequence $PrivKey_0$, thus $\rho_k^0 = |b_k^0\alpha\rangle\langle b_k^0\alpha|$ where α is a real positive amplitude. A sequence of such states is called a quantum signature. She sends a copy of the quantum signatures to each of Bob and Charlie each, informing them that they correspond to message $m = 0$. Alice then does analogously for the message $m = 1$. The individual state ρ_k^m we refer to as the k^{th} quantum signature element state for message m.

3. Bob and Charlie send their copies of the sequences $QuantSig_0$ and $QuantSig_1$ through the multiport. The exit nullports on Bob's and Charlie's side of the multiport are equipped with photon detectors and the total number of photon events here will serve to disable certain types of forging attacks, but are not crucial for security against message repudiation.

4. As the elements of the quantum signatures exit the multiport Bob and Charlie perform optimal unambiguous state discrimination (USD) on these states, storing the outcomes which are in $\{-1, 1, \text{“?”}\}$, corresponding to the states $|-\alpha\rangle$, α and the ambiguous outcome. We define the USD measurement, along with the meaning of the outcomes in the main body of text. These *classical outcomes* are stored in classical memory. For the sequences of classical outcomes they get, Bob and Charlie note to which message $m = 0$ or $m = 1$ they correspond to.

5. To sign a single bit, say $m = 0$ with Bob, Alice announces the message m and the corresponding private key to Bob. Thus, she sends the pair $(0, PrivKey_0)$ over an untrusted channel. To authenticate the signature, Bob checks whether the phases Alice declared match the phases he measured at the signature elements for which he received the unambiguous outcome. He authenticated if the number of mismatches is below $s_a L$ for an **authenication threshold** s_a which is zero in the ideal case.

6. To prove to Charlie that he received the message $m = 0$ from Alice, Bob forwards to Charlie the pair $(0, PrivKey_0)$ he received from Alice. Charlie then performs an analogous procedure to Bob, and he verifies the message (*i.e. the message passes verification*) if the number of mismatches is below $s_v L$ where s_v is called the **verification threshold**, with $0 < s_a < s_v < 1$.

The protocol is aborted if any of the thresholds are breached.

Passive attacks We first consider passive attacks, meaning Bob did not act maliciously during the distribution phase *i.e.* within the multiport.

For a signature of lengths L, on average, Charlie will have obtained pL unambiguous measurement outcomes. This probability p for the case of two possible states $|\alpha\rangle$ and $|-\alpha\rangle$ equals $1-e^{-\alpha^2}$, assuming α is real and positive.

Since the phase signs are chosen by independent random variables, the optimal strategy of Bob is to perform individual measurements which minimize the probability of him making an incorrect guess [5]. By definition this is the minimum-error measurement, and for the two coherent states occuring with equal probability, the success probability of the minimum-error measurement is $p_{min} = \frac{1}{2}\left(\sqrt{1-e^{-2a^2}}+1\right)$. The probability Bob guesses the correct phase for all of the pL states is then

$$p_{forgery}^{passive} = p_{min}^{pL},$$

which is decaying exponentially quickly in L. So we have satisfied the definition of security against forging for the passive attack.

Active attacks In complete analogy to the analysis we performed for the experiment in the previous chapter, by keeping track of the null-port multiport counts we can again guarantee the following, at least for the case separable attacks. If Bob's malevolent activity did not cause a violation of the multiport count threshold then quantum signatures Charlie receives are per element in terms of the trace distance $\delta-$close to unperturbed quantum signatures. This value δ can be controlled by the key length and multiport thresholds, and be made arbitrarily small. For the details of these types of attacks and the security analysis, please refer to Section 8.3.4. Because of the operational characterisation of the trace distance (elaborated in Section 8.4.2), the measurement outcomes Charlie obtains cannot differ substantially from what he would obtain if malicious Bob employed a passive strategy. More precisely we would obtain an expression for Bob's cheating probability of the following form [6]:

$$p_{forgery}^{active,separable} \leq (p_{min}+\delta)^{pL} \tag{9.3}$$

where δ is the trace distance between the quantum signatures Charlie measures in the passive or active separable strategy employed by malicious Bob.

Unfortunately, the type of reasoning we employed here does not trivially extend to coherent active strategies, but the same argument used in the analysis of active coherent forging attacks in Section 8.3.5 can be applied here as well.

[5] To see that coherent measurements cannot help we refer the reader to the analysis in Section 8.3.3 where we proved a more general claim.

[6] The expression we present here may not be exact, but the asymptotic behaviour is correct.

9.2.1.2 Security against repudiation

For the ideal case, with an ideal multiport, the state of the system going to Bob and Charlie, when the multiport null-systems have been traced out are symmetric under swap. This observation was at the crux of the proof of security against repudiation in the original coherent state QDS protocol, and will also guarantee that a malevolent Alice cannot make Bob and Charlie disagree on the validity of her message in this modified protocol.

Arguments are analogous to the ones presented in sections 8.2.1 and 8.2.2. Honest Bob and Charlie perform identical USD measurements on the systems which exit the multiport signal arms. If these states have identical reduced matrices, then the effects, the induced probability distributions of measurement outcomes, that Bob and Charlie see are identical.

But this again implies that Bob sees a mismatch for a particular element state he measured, and obtained an unambiguous outcome, relative to what Alice declares as her private key with the same probability as Charlie in the same setting. Again, by slightly raising the verification threshold compared to the authentication threshold, we can make the probability an authenticated message fails to pass the verification test arbitrarily small. This, of course does not constitute a formal proof, but the same ideas used in the full analysis in the Chapter 8 should apply in this case as well. This includes the setting when coherent strategies by Alice and experimental imperfections are taken into account as well.

9.2.2 Experimental realizability of the QDS protocol with no quantum memory

The simple protocol we have described can easily be generalized in many ways, but the variant we propose is very convenient for implementation. The only difference between the proposal in [3] and this modified protocol is that the quantum signature elements are measured during distribution, and only classical outcomes are stored, and the measurement performed as proposed is the optimal USD measurement. Optimal unambiguous discrimination of the two coherent states $|\alpha\rangle$ and $|-\alpha\rangle$ is easily performed in linear optics: one simply prepares the state $|\alpha\rangle$ and combines it with the unknown state on a balanced 50:50 beamsplitter and observes photon counts on the two output arms. When no imperfections are present, only one of the arms will contain light, depending on what the input state was, and the detection of the photon conclusively discerns the input state. However, the outputs will be coherent states again, which contain a vacuum component, so it is possible, and in the case of small amplitudes very likely, that neither of the detectors fire. This constitutes the inconclusive result. In the case no imperfections were present this is also the optimal measurement.

Thus, this proposal is experimentally friendly.

However, if it would be convenient to use not just two states, but rather many, then optimal unambiguous discrimination probably can no longer be performed using linear optics alone.

But, in principle, this protocol may work even if unambiguous discrimination is replaced with a minimum-error measurement, which may be easier to implement than unambiguous discrimina-

tion in optics. In this case, the required thresholds which guarantee security cannot be trivially established, even in the ideal case, and the security analysis may be more involved. For this reason, we chose to present this protocol based on USD measurements.

9.2.3 Quantum state comparison

As noted, any "quantised" version of the Lamport's scheme (see Section 6.2) will have to deal with the issue of quantum state comparison, to ensure security against repudiation. In our analysis of the coherent state-based protocol, which uses the multiport, the property that the multiport can be used compare quantum states entering it was not relevant for security against repudiation, but rather for the security against active forging attacks only. What we did extensively use is the property that the states Bob and Charlie store, or immediately measure in the variant of the protocol proposed in the previous section, are symmetric under the swap of the systems of Bob and Charlie. The multiport, although a passive device, could be very difficult to build in practice if large distance between the receiving parties are required, and in any practical scenario, this would be the case. Ideally, we would be able to somehow replace the multiport, which realizes a global operation on the systems of all recipients, with local operations, equipped perhaps with classical communication. This is unlikely to be achievable for general states. However, there may be local method which allows Bob and Charlie to be certain that the states they have received are arbitrarily close to each other by using their knowledge of what they were meant to receive – by checking if Alice is being honest through classical communication which includes her as well.

We will now present a protocol which intuitively may achieve the goal presented, and then discuss how its validity could be proven in certain cases. Our goal is to produce two quantum states, one with Bob and one with Charlie, for which we have guarantees that they are close in terms of a relevant distance measure (ideally trace distance) and yet, Bob and Charlie have no additional information about it (guaranteeing security against forging). If one such state can be generated with just direct quantum channels from Alice to the recipients (and classical channels connecting all parties), then by using the procedure many times we can run the entire QDS protocol. A sketch of this scheme is given in Protocol 16.

Intuitively it seems in this protocol that since Alice does not know in advance which state will be selected as the unmeasured state by Charlie and Bob, if the protocol is not aborted in the limit $M \to \infty$, then Alice was honest. However, making a quantitative statement about the j^{th} states Bob and Charlie keep for a finite L, when Alice is not restricted in any way seems difficult. Nonetheless, we may be able to use "heavy machinery" developed for QKD, if we slightly tweak the protocol and use non-orthogonal qubit states instead of coherent states. Since qubit states are of finite dimension, we conjecture a quantitative statement could be made by employing a variant of the quantum de Finetti theorem [114] (which does not hold in all generality for infinite dimensional systems).

To explain how this may work we shall use a bit more detailed notation. After Alice send away

Protocol 16 Quantum states comparison through local operations and classical communication

1. Alice then generates two copies of a sequence of *e.g.* coherent states $QuantSig_0 = (\rho_1^0, \ldots, \rho_L^0)$ with the coherent phase of the sign matching the signs in the sequence $PrivKey_0$, which is again a string of signs b_k^0, thus $\rho_k^0 = |b_k^0\alpha\rangle\langle b_k^0\alpha|$ where α is a real positive amplitude.

2. After they receive the states from Alice, Bob and Charlie store the states in *short term quantum memory*.

3. Following this, they each choose random integers int_{Bob} and $int_{Charlie}$ chosen from 0 to $L-1$, exchange these integers over authenticated classical channels, but also send them to Alice.

4. Alice computes $j = int_{Bob} + int_{Charlie} - 1 \mod L$ and reveals all the choices of phases except for the j^{th} element.

5. Bob and Charlie compare the classical strings they receive from Alice, using their authenticated channel to make sure Alice send them identical strings and that j indeed is the mod N sum of their random integers.

6. Bob and Charlie verify whether the phases Alice declared match with the states they received for all but the j^{th} state, by performing the projective measurement characterised by POVM elements $\Pi_{ok}^k = |b_{k\ declared}^0\alpha\rangle\langle b_{k\ declared}^0\alpha|$ and $\Pi_{abort}^k = \mathbb{1} - \Pi_{ok}^k$.

7. If any measurement result corresponds to the Π_{abort}^k measurement operator the protocol is aborted.

8. The j^{th} state is kept as a quantum signature element.

her stated to Bob and Charlie the state they share is a $2M+1$ partite state:

$$\pi_{A,B_1,C_1,...B_L,C_L}.$$

The system A is controlled by Alice, but for simplicity let us assume Alice will not use her system in order to cheat. In the analysis against repudiation in Section 8.2.2 we have shown that this can be done without the loss of generality for the case when the multiport is employed. Whether this assumption can safely be made in the case we consider now, we will address later. Thus, Alice's system we trace out leaving us with

$$\pi_{B_1,C_1,...B_L,C_L}. \tag{9.3}$$

In this protocol Bob and Charlie by using authenticated and private channels effectively decide which state they will leave unmeasured, but can also agree on the order in which the measurements will be performed. Alice can see this as follows: she sends away the state $\pi_{B_1,C_1,...B_LC_L}$. Bob and Charlie randomly permute all their pair-wise matched systems and will measure all but the last state always. Thus prior to measurement the state is a mixture of states of the form

$$1/M! \sum_{\sigma \in S_L} \pi_{B_{\sigma(1)},C_{\sigma(1)},B_{\sigma(2)},C_{\sigma(2)}...B_{\sigma(M)}C_{\sigma(M)}} \tag{9.4}$$

summed over all permutations σ in the symmetric group S_L. This trick we may call "symmetrization", and a similar method was used in QKD. There, through symmetrization, one of the proofs of security of QKD against general attacks was produced. The main idea in these types of security proofs for QKD was reducing security against general attacks to security against collective attacks through the use of a version of the quantum de Finetti Theorem [114].

The state in the equation above is clearly invariant under permutations of the signature elements. But then, by the quantum de Finnetti theorem, by discarding a moderate fraction of the systems, and Bob and Charlie may agree on which will get discarded using their private authenticated channels, the remaining state will in terms of the trace distance be (arbitrarily) close to a product state of identical copies of quantum element states shared by Bob and Charlie. By performing the measurements according to Alice's declaration, Bob and Charlie effectively perform a type of state tomography – they are measuring many copies of the identical state [7]. Possibly, now a statement about how close the resulting state j is to the state honest Alice would have sent can be made. The direct application of the basic de Finetti type approaches to this problem will not yield efficient scaling of the realized security. However, it is possible that more advanced techniques used in QKD which allow exponential improvement on the basic result, such as the exponential variant of the de Finetti theorem [114] or the so-called postselection techniques [115], might offer substantial improvement in this situation as well.

In the arguments above we have assumed Alice does not use the option of keeping a state entangled to the systems she sends to Bob and Charlie for future use. In the original QDS scheme

[7] Actually, they are identical with respect to a varying frame of reference, where the variation is defined by Alice's declaration of the phases.

we analysed this was justified as no classical communication existed between Alice on one side, and Charlie and Bob on the other. However in this protocol, Bob and Charlie inform Alice of their choice of j. Whether Alice can, in this setting, use entanglement to somehow improve her cheating capabilities, and formal proofs (or disproofs!) of the conjectures we have made in this section we leave as open problems.

Experimental realizability of quantum state comparison without global interactions As we have mentioned, the proposed protocol could potentially be proven to work using existing techniques, provided the systems we use are qubits rather than coherent states. In general, working with coherent states is easier for the experimentalist. Many interesting measurements, like unambiguous state discrimination of two states, can be performed just using linear optics and the states involved are generated by cheap and yet high quality laser sources. Working in the single photon regime, which is one way of realizing the qubit setting may be more difficult but should not be impossible. Alternatively, a proof of a theorem could potentially be produced which quantifies the "quality" of the state Bob and Charlie end up with in the protocol we describe, however at this point the author cannot estimate how difficult this could be, and if it is at all possible.

In the scheme we present, we assume Bob and Charlie have an amount of short-term quantum memory. In particular, they have to keep their pulses stored long enough to generate their random integers, this can be done off-line, send them to Alice, and await her responses. In reality, the slowest process here is the travel time of information from Bob and Charlie to Alice and back, if we assume Alice is reasonably far from them. But again, this memory, unlike the memory required for QDS with quantum memory, is of perhaps a reasonable duration [8]. It is also conceivable that event this short-term quantum memory can be completely avoided by cleverly choosing the timings of measurement performed in run-time. The two proposals for the modification of the existent QDS scheme can be used in conjunction. The quantum states denoted j in the protocol above can immediately be measured using unambiguous discrimination and stored, as the classical information Bob and Charlie will use as their "classical signatures". Note that the amount of classical information Bob and Charlie need to store is actually less, in terms of bits vs. qumodes, even when the distinction between the two is ignored, than in the protocol with quantum memory. They only need to store the unambiguous outcomes and the ordinal number of the pulse when that outcome occurred.

This advantage is perhaps analogous to the setting in QKD realizations, where the events when no photons were detected at all are simply gated out and do not contribute to the raw keys stored and further manipulated. In QDS with quantum memory, all the modes need storing, since we cannot know which ones will result in no detection events at any of the proton detectors in the authentication/verification stage.. Presuming all the proposed claims can actually be proven to work formally and rigorously, and provided the security holds once realistic imperfections are

[8] A simple, although not very practical idea how to implement the required quantum memory would be for Bob and Charlie to simply have a spool of low loss, low noise optical fibre in their labs, approximately twice the length of the distance from Alice to them. However, this would have issues with stability due to mechanical and thermal fluctuations in path length.

considered, what we may have is a QDS scheme which in principle can be realized with the same equipment used even in commercial QKD systems.

So what is quantum in QDS anyway? One legitimate question which can be asked is the following: what is the inherently "quantum" property which QDS relies on? In the initial proposal the answer to this would possibly be the no cloning property of quantum states. However, in the discussions of this chapter, we have suggested that most of the quantum-looking components, like quantum memory and the multiport can be dealt away with. What remains, and is fundamental for the schemes to work, and cannot be achieved classically are *noisy classical channels* with *guaranteed levels of independent noise sources*. If Alice had classical channels to Bob and Charlie which are guaranteed to transmit a fixed percentage of classical bits, where which bits get transmitted is determined at random, and independently for Bob and Charlie, Alice would not need to resort to quantum states. Such considerations have also been considered for the problem of sharing secret keys. If upper bounds on the transmittance of a special type of a classical channel can be guaranteed, then secure key expansion, of the type achieved in QKD can be performed without quantumness as well. Such observations have been explored before in similar contexts [116, 117]. Since quantum effects guaranteeing independent noise sources seem to be the key thing exploited in both QDS and QKD, this opens up the natural question of how distinct QDS and QKD really are. However, such considerations only become pertinent if all assumptions usually coming with public versus private key cryptography are dropped, as we will briefly comment in the following section.

In the following section, we take yet another step back from the details of the protocol, and investigate the potential place QDS may take in the future.

9.3 Quantum digital signatures and the big picture

The classical and quantum cryptography communities unfortunately seem to be not completely synchronized. In the theory classical cryptography, as we have mentioned in the introduction to Part 2 of this thesis, a key distinction exists when one considers private of public-key cryptosystems. Depending which setting one works in, different assumptions tend to be allowed. The one issue with QDS is that it is unclear whether it can justly be considered a public-key proposal, since the "public keys" are quantum states. If this assumption is justified, then we can always assume that the initial point of any security analysis of QDS begins with the recipients already having the quantum keys Alice wants them to have. Alternatively, if we do not have this assumption, then quantum channels need to be authenticated, in an information-theoretically secure way, and we as immediately in a public-key setting. Practical cryptography tries to deal with the difficult issues of safely distributing keys, both private and public, but this always necessarily involves frowned upon, but unavoidable entities like trusted couriers, and trusted centres. Quantum key distribution is a well-established field, which clearly fits in the private-key setting. For QDS, the situation is less clear. Here, we do not presume to argue whether QDS should or should not be considered a public-key scheme from a theory point of view. But, this necessarily forces

us to drop as many problematic assumptions as possible, which includes considering dropping authentic distribution of public keys. This seemingly pushes QDS into the private-key domain, in which case, by current definitions, this is not a digital signatures scheme at all. As scientists in the quantum information processing field, we are forced to think about the actual physical reality, which seems to suggest we should drop as many assumptions as possible, even at the price of pushing QDS into the private-key setting. We do concede this may be very delicate, however, at least throughout this discussion chapter of the thesis, we shall, perhaps stubbornly, throw caution to the wind, and think od QDS in terms of its functionality, and not a protocol of public or private key species. We do acknowledge that this constitutes a dangerous mixing of the theory of cryptography, and the practice of cryptography [103].

One of the frequent criticisms one hears concerning the QDS through discussions with classical cryptographers is that the functionality it achieves is only seemingly related to real-life classical digital signatures. Another issue coming from experts in classical cryptography is that QDS seems overly resource expensive. Other classes of issues raised, coming form cryptographers standing firmly in both the classical and quantum realm, point out the resources required during the quantum signatures distribution phase which we simply assume are given, but do not account for the overall cost of running a QDS scheme. In this section we address those issues, and in the end address additional remarks made by other authors in published materials, which are relevant for the future of QDS.

We begin by investigating the pros and cons of QDS relative to standard digital signatures, as functionalities.

According to the authors in [118] there are two significant disadvantages of the original Gottesman-Chuang scheme (QDS) compared to conventional digital signatures (cDS) schemes used in practice. We combine their objections with additional concerns we raise.

Accessibility of quantum signatures (public keys) In cDS the distribution of public keys is of low cost, because any participant can easily extract other's public keys from a public directory. Moreover, the number of participants does not directly influence the overall (computational) security of the protocol. On the other hand in QDS, the quantum signatures – public keys – need to be distributed over an authenticated line, either by a trusted centre or Alice. Furthermore, in QDS the number of participants has to be limited, and accounted for in the required lengths of the signatures. Most likely one should agree that, in practice, simply assuming a trusted centre, equipped with authenticated channels to all participants is not a valid solution. Alice could then use the trusted centre to do all the transactions for her using the pre-existing authenticated channels. The only potential solution seems to be that Alice herself to have *authenticated quantum channels* to all participants. We will shortly investigate what the cost of this would be, but also suggest a potential solution which may ameliorate this requirement.

The limitation of the allowed number of recipients being could, in principle be ameliorated by brute-force. For instance, Alice could establish quantum keys with the rest of the world, and given keys of sufficient length, still have security. Note that the security of our proposed system

scales exponentially with the signature length, whereas the accessible information is only linearly dependent on the number of copies. However, it may not be completely fair to compare the worst case scenario for QDS with many recipients where *all but one* of the recipients collaborate to forge. One could possibly argue that if *all the possible users* of a cDS scheme were to combine their computational powers in order to cheat on the one remaining participant, the currently used cDS schemes would not provide security guaranteed for an acceptable time-frame. As we noted, most public-key cDS schemes can be broken, given sufficient time and raw computing powers by brute force. However, such an argument should not be taken too seriously. The next concern is perhaps a greater issue.

Durability of public keys In most cDS systems, the same public keys can be re-used in signing different messages (perhaps only simple adaptations need to be performed). However, in QDS (and Lamport's scheme) the keys *cannot* be re-used. This is an important issue as it limits the sizes of messages one may want to send over authenticated channels, or alternatively, requires huge storing resources to hold sufficient number of public keys for all messages we *ever* may want to send.

Recently, Ioannou, Lawrence and Mosca in [119] showed that unconditionally secure and reusable public-key authentication is possible in principle with pure-state public keys. What they achieve is a limited re-suability of the keys, *i.e.* the keys can be re-used a prescribed number of times. Further advances on this initial result may still result in re-usable quantum public key cryptography. But, the re-usability will still be limited.

This advantage of cDS is serious, but, we have not yet weighed in the main advantage of QDS - information-theoretic security. For many applications, the desired levels of security may be not overly extravagant. However, from time-to-time the requirement of high levels of security may trump any less secure, but more cost-effective solution. For instance short term limited security may be very problematic if we wished to run cryptographic protocols on highly sensitive data, such as crucial medical or financial data or government secrets. Such data often has to remain confidential for significant fractions of centuries. This problem certainly applies for the secrecy of messages, but could potentially be an issue in authentication-type tasks as well. In this setting we suggest that in the modern world, where quantum technologies are becoming increasingly accessible, *both* types of solutions for the problem of signing messages should be used, depending on the required level of secrecy. Estimating the realistic cost of any activity – such as choosing to use a particular cryptographic protocol – should also include the *risks* which accompany any action, not just the resources required [9]. How does one evaluate the cost of limited security in general? While this question may not have a sensible unique answer, simply ignoring the fact that security is not unconditional, or at least everlasting, should not be satisfactory.

If one is willing to accept that it may make sense to use non-reusable cryptosystems for signatures, when they obtain information-theoretical security because of it, then we may as well see

[9] To push the issue to extremes, what if supercomputational Aliens attack Earth? Well, one only hopes such complexity-rules-violating aliens are well behaved enough not to violate standard quantum mechanics or no-signalling...

what the other costs of such a system are, and whether they can be reduced.

Authenticated quantum channels in QDS As we have noted, in the three party QDS we assume Bob and Charlie share *classical* authenticated channels, which prevents Alice from repudiating her message by disrupting their communication. This, arguably, is not a big issue, because as we have explained, even large messages can be authenticated using a very short pre-shared key. However, having authenticated quantum channels is a different story. Not much is known on authenticating continuous variables communication, *i.e.* communicating systems with infinite degrees of freedom which we use in our proposals. But, even for qubit, systems it is known that the size of a pre-shared key has to be of the length of the order of the number of qubits we wish to send [17]. In the case of quantum channels, authentication is as difficult as ensuring privacy. Intuitively this should not be overly surprising, as we know that eavesdropping in the quantum realm is an active event – it is a measurement, and measurements perturb the state. Thus, if our protocol ensures protection against tampering – perturbing the state – it may well provide security against eavesdropping. And the price for one-time padding a qubit is known to be two classical bits. In the QDS protocol, and particularly in our realization, the number of quantum states we have to send, the length of quantum signatures is very large, even for a single bit. Thus, a large number of secret shared bits is required. So large in fact, that one should first use QKD, perhaps in run time, to generate them. Thus QDS powered by QKD could solve our problem, but the communication cost radically increases.

However, it may be possible to kill two birds with one stone. In this chapter we have proposed a modification of the QDS scheme, which requires no multiport, or any other type of quantum state comparison / symmetrization device (see Section 9.2.3). This proposal requires a certain amount of authenticated classical communication between Bob and Charlie, and quite a bit of classical communication from Alice to Bob and Charlie, during the phase in which she reveals most of her private key. Since this is classical communication only, authentication of these messages is relatively inexpensive. However, if the classical information arrived correctly to, say Bob, describing the states of the quantum pulses Alice sent, by measuring all but one of them (the "j^{th}" state, in our description), Bob verifies whether the pulses are not only the same as Bob's, but also whether he received what Alice sent. This may be a viable a method for authenticating a quantum state, which we get for free. If this idea works, it still remains to be evaluated whether it outperforms QDS powered by QKD in terms of overall communication and final memory cost. It is very likely that our proposed method actually uses aspects of QKD, or realizes them, implicitly. It would not be surprising that QDS and QKD are actually much more closely related than previously thought.

It is known that unconditionally secure key exchange is also possible without the use of quantum mechanics, if resources with a guaranteed level of noise are available [116]. The advantage of QKD from this perspective is that, in a sense, here "noise" is guaranteed by quantum mechanics and it is quantifiable and guaranteed, as we have mentioned earlier.

Could there then exist *fully classical* variants of digital signatures, which allow a fixed number of participants and achieve unconditional security? Probably yes.

Unconditionally secure classical digital signatures The protocol is explained in Protocol 17, where we show how half a bit may be signed [10]. For a full bit, the protocol is run twice, and the messages exchanged are denoted whether they correspond to message "0" (in, say, the first run) or "1" (second run).

Protocol 17 Classical unconditionally secure digital signatures – D. Unruh

1. Bob randomly chooses a key key_B, and a hash function h_B from a set of universal hash functions. He sends key_B to Alice, and $(h_B, h_B(key_B))$ to Charlie.

2. Charlie does the same as Bob.

3. To sign half a bit with, say Charlie, she sends key_B^{Alice} (Bob's key) to Charlie. For honest Alice $key_B^{Alice} = key_B$

4. To authenticate the bit, Charlie applies the hash function h_B he received from Bob to the key key_B^{Alice} he got from Alice, and checks if it matches $h_B(key_B)$ he received from Bob. He authenticates if and only if it does.

5. To verify the message with Charlie, he forwards key_B^{Alice} to Bob, who verifies if $key_B^{Alice} = key_B$.

   All the channels are private and authenticated. The hash functions go from a large set from a smaller one. The sizes are the security parameters.

What follows is a few-line mock security analysis. Bob cannot pretend he is Alice, as all parties are assumed to share authenticated channels (corrupted Bob cannot pass authentication). He cannot pretend he received a message from Alice, which he did not, as he cannot generate the valid pre-image key_B from h_B and $h_B(key_B)$ except with negligible probability due to the properties of these (non-invertible) hash functions. Alice cannot repudiate a message, because she does not know the hash function h_B, so what protects Bob and Charlie are the no-collision properties of universal hash functions.

While this protocol seems to achieve unconditional security, the amount of private and authenticated communication is substantially larger than in QDS (in particular QDS requires no private communication), and thus requires an exorbitant amount of pre–shared keys. One can then upgrade the proposed protocol by growing the long keys required for private channels using QKD. Now, certain quantum resources are required, but the infrastructure and technology for those exist right now. However, should our proposed no quantum memory and no multiport QDS scheme be proven secure, it is unclear which solution to unconditionally secure digital signatures is more efficient. This question warrants further research.

Other proposals for information-theoretically secure digital signatures schemes which are completely classical exist [120, 121]. Similarly to QDS in these schemes the total number of users is limited (pre-defined), and all the parties involved need to store a substantial amount of classical

[10] Towards the end of this project the author had the opportunity of discussing with Prof. Dominique Unruh, who proposed the following scheme, which achieves the same functionality as QDS and unconditional security, and requires no quantum resources.

information to ensure security. However, unlike in the schemes we have presented so far, the proposals in [120, 121] assume the existence of a trusted authority for the protocol to work. At this point it is unclear whether the trusted authority could be replaced by pre-shared secret keys. In the likely event that this is possible, the true relationships between all the classical schemes and quantum schemes for digital signatures will depend on the overall efficiency – cost per signed bit – in each scheme. Thus far, the unconditionally secure schemes for digital signatures have received modest attention from the classical cryptographic community. At the present state of affairs, the efficiency and lack of requirement of pre-shared keys in computationally-secure digital signatures schemes, seem to outweigh the benefit of information-theoretical security. If the moment should occur when machines which easily violate some of the computational assumptions are built (for instance quantum computers, or specialized devices designed for a fixed instance size), we may see a shift in this trend.

Chapter 10

Side results: *some properties of symmetric sets of states and applications*

Here we investigate the properties of symmetric sets of quantum states, which appear both in the presented UBQC and QDS protocols. We show they admit an elegant theory characterising the possible transforms between such states. We demonstrate two relevant applications of the theory.

In both protocols of UBQC and QDS this thesis has addressed, the security properties for the users were guaranteed by encoding classical information in particular types of quantum states. In UBQC, encryption angles parametrized the qubit states the client sent to the server, which allowed for a correct, yet fully concealed computation. In the AKLT UBQC case, more elaborate "'Dango" states were used for the same purpose. In QDS, the "quantum keys" which ensured security against forgery and repudiation were coherent states, with complex phases encoding the secret information fully known to the honest sender alone.

All the quantum states which appeared in both protocols (which include the classical messages as well) comprised *symmetric sets of quantum states*. This symmetricity. which we formally define presently, was crucial in the security proofs. In this section we exploit this symmetricity to devise a comparatively simple theory which characterises properties of transformations between symmetric sets of states. As applications, we present results concerning transforms which convert phase encoded coherent states (used in QDS) into relative-phase encoded qubits (used in UBQC), and also results on truly perfect amplification of phase-encoded coherent pulses, which can be achieved probabilistically under realistic assumptions.

While the research presented in this chapter played a role in the security proofs we presented throughout this thesis, the results we give are more general. Thus we present them in a more general context, and this chapter can be read independently from the rest of the thesis.

10.1 Transformations between symmetric sets of quantum states

10.1.1 Introduction

Quantum information theory promises new and exciting ways to process information. Often the advantage a quantum protocol gives over a classical procedure lies in the fact that quantum

states may be non-orthogonal. Classical information may be encoded in non-orthogonal quantum states, as is the case for example in quantum key distribution. The classical information then cannot be fully extracted from the quantum state alone. Many physical systems are candidates for the realization of quantum processing, and often they perform well at distinct tasks. Therefore, future quantum devices may well be hybrid systems with interfaces linking the different parts, just as our classical information processing devices are today. When classical information is encoded in a set of quantum states, the information encodings of one system may be incompatible with the encodings of another in a sense which has no classical analogue: the 'source' states may not be related to the corresponding 'target' states by a fixed unitary transformation. This occurs, for instance, when we consider transforms between states of systems of distinct dimensionalities such as typical qubit states and coherent or sqeezed states of continuous-variable systems. When transferring information from one system to another where the encodings are incompatible in this sense, we must then either accept errors or resort to probabilistic scenarios where information may be lost, or leaked. This is important from an information-theoretic and cryptographic perspective. Information is no longer perfectly controlled by the emitting party.

Transformations that take a 'source' set of quantum states to a 'target' set of quantum states, where the states in the two sets are not pairwise related by a single unitary transform, also have other applications. State-dependent quantum cloning is one example [122]. Another related and well-studied family of such transforms solve the problem of amplifying coherent light, while keeping the coherent phase unaltered [123, 124, 125, 126, 127]. This problem is very important in classical and quantum communication tasks over larger distances, and is usually resolved by generating approximations of amplified coherent states. Optimal measurements for distinguishing between quantum states can also be seen as transforms taking some set of quantum states to mutually orthogonal states, followed by a measurement to distinguish the latter from each other. For so-called minimum-error measurements, pioneered by Holevo and Helstrom, the transforms are allowed to err, i.e. the declared output need not always be correct. Another tradition requires correctness but allows for result which declares that the transform (measurement) has failed, following the works of Ivanović, Dieks and Peres [128, 129, 130]. Such measurements are then called unambiguous [131, 132].

Here we focus on unambiguous transforms taking pure states to pure states. For this setting there exists a convenient framework based on the structures of the Gram matrices of 'source' and 'target' states, developed by Chefles, Jozsa and Winter [133, 134], which we will briefly present. In these works, the sets of source and targets states are general, and finding transforms for given sets of source and target states is complicated. However, it is known that for the problem of distinguishing quantum states which comprise a symmetric set, a simpler treatment is possible [131, 113, 135]. As we will show, this restriction simplifies the theory for general probabilistic transforms as well. As an application of the theory we develop, we study the properties of converting a set of coherent states to qubit states. This is an important example of an 'interspecies' transform, as these two types of encodings frequently appear in quantum information processing tasks.

10.1.2 Preliminaries

Our problem of interest is stated as follows: given two sets of pure states A and B (called 'source' and 'target' sets, respectively) of finite size N,

$$A = \{|a_i\rangle\}_{i=1}^{N};\ B = \{|b_i\rangle\}_{i=1}^{N},$$

what are the properties of a transform $\mathcal{T}$, allowed by quantum mechanics, which performs $\mathcal{T}(|a_i\rangle) = |b_i\rangle$ for all i perfectly with a certain probability? The transform can fail to produce the desired output state, or succeed, and these two possible outcomes are reported, i.e. the transform is heralded.

In the most general case, the success probabilities may depend on which source state we start from. In the Lemma which follows, $\circ$ denotes the Hadamard (Shur, point-wise) matrix product, and the Gram matrix of the set of kets (or more generally vectors) $A = \{|a_i\rangle\}_{i=0}^{N-1}$ is given by

$$G_A = \left[\langle a_p|a_q\rangle\right]_{p,q},\ p, q = 0, \ldots, N-1.$$

The necessary and sufficient conditions for a matrix M to be a Gram matrix of normalized kets (states) are that $i)$ M is a positive-semidefinite matrix, and $ii)$ M has unity across the main diagonal.

We then have the following statement:
Lemma 30. *There exists a probabilistic transform taking each state $|a_i\rangle$ in A to the state $|b_i\rangle$ in B, succeeding with the (non-zero) probabilities p_i, for $i = 1, \ldots, N$, iff there exist Gram matrices of kets Π^s and Π^f such that the equality*

$$G_A = P^s \circ \Pi^s \circ G_B + P^f \circ \Pi^f \tag{10.-1}$$

holds, where

$$P^s = \left[\sqrt{p_i p_j}\right]_{i,j} \text{ and } P^f = \left[\sqrt{(1-p_i)(1-p_j)}\right]_{i,j}, \tag{10.0}$$

and G_A and G_B are the Gram matrices of sets A and B.

This is a special case of the Theorem 3 in [133]. Such a quantum transform can be equivalently viewed, in the spirit of the Stinespring dilation, as a unitary transform acting on an augmented Hilbert space,

$$U|a_i\rangle|0\rangle|0\rangle = \sqrt{p_i}|b_i\rangle|\psi_i\rangle|0\rangle + \sqrt{1-p_i}|Fail\rangle|\phi_i\rangle|1\rangle \text{ for all } i, \tag{10.1}$$

where we learn whether the transform has succeeded or failed by measuring the third 'indicator' register on the right-hand side of the expression. One may show that the matrices Π^s and Π^f in expression (10.1.2) are the Gram matrices of the sets of kets $\{|\psi_i\rangle\}_i$ and $\{|\phi_i\rangle\}_i$, respectively. If

the transform in equation (10.1) succeeds, then the output registers contain the target state $|b_i\rangle$ but also a residual state $|\psi_i\rangle$ which may be correlated with the input state. From an information-theoretic perspective, this residual state may be seen as a leak of information, hence we call the set of states $\{|\psi_i\rangle\}_i$ the *leak*. If the states $|\psi_i\rangle$ are not correlated with the input state, which happens if and only if $|\psi_i\rangle = |\psi_j\rangle$ for all i and j, then the transform is called *leakless*. Analogously, in case the transform fails, a fixed fail state is produced along with a residual state $|\phi_i\rangle$ is produced. The residual state may be correlated with the input state, and may be used to subsequently attempt to reconstruct the desired outcome. For this reason we call the set of states $\{|\phi_i\rangle\}_i$ the *the redundancy*. If all the states in the redundancy are identical, only then is the residual state uncorrelated to the input state, and the transform is called *redundancy-free*.

If the success probabilities above do not depend on the source state ($p_i = p_j$ for all i, j), we call the transform *uniform*. For this case the notion of the optimal transform can be naturally defined: a uniform probabilistic transform is optimal, if no other transform succeeds with a strictly grater probability.

Deterministic and unitary transforms are easily seen to be special cases of probabilistic transforms. For a deterministic transform, it holds that $p_i = 1$, in which case the criterion reads $G_A = \Pi^s \circ G_B$ (for some Gram matrix of states Π^s). For a unitary transform the complex matrix Π^s is an outer product of a vector, containing roots of unity, with itself (c.f. [133]) [1]. Throughout this section, with G_S we will denote the Gram matrix of the set of states S, and with λ_{G_S} a vector comprising the eigenvalues of the matrix G_S. With I we denote the identity matrix, and with $\mathbf{1}$ we denote the matrix with unity at each entry.

10.1.2.1 Example: uniform unambiguous discrimination of pure states

Unambiguous discrimination of states (UDS) identifies the input state from a pre-defined set of states, error free, but allows a 'failure' option. It is equivalent to a probabilistic transform for which the states $|b_i\rangle$ are mutually orthogonal. The criterion for the existence of such a transform is given by Lemma 30. Since the Gram matrix of orthogonal states is the identity, and the Hadamard product of the identity and Gram matrix of states is the identity again, for the special case where the success probability p is independent of the source state (uniform UDS), the existence condition simplifies to the inequality

$$G_A - pI \geq 0, \tag{10.2}$$

meaning that the matrix $G_A - pI$ is positive-semidefinite. Since unitary basis change preserves operator positivity, and G_A is positive-semidefinite, hence diagonalizable in an orthonormal basis, this implies and is implied by

$$p \leq \min \lambda_{G_A}, \tag{10.3}$$

[1] This freedom in the complex phases reflects the fact that kets in general contain information about the physically irrelevant complex phase.

where $\min \lambda_{G_A}$ denotes the smallest eigenvalue of the matrix G_A. From this condition we easily capture a known result: the optimal success probability of UDS is equal to the smallest eigenvalue of the Gram matrix G_A [2]. A consequence of this is another famous result: a set of states may be unambiguously discriminated if and only if that set of states is linearly independent. The latter is clear as the spectrum of G_A contains a zero element if and only if the set A is linearly dependent.

From the fact that unambiguous state discrimination is possible iff a set of states if linearly independent, it is easy to see that if a uniform probabilistic transform $\mathcal{T}$ is optimal, then the redundancy is a linearly dependent set of states. To prove this, assume that a uniform probabilistic transform $\mathcal{T}$ succeeds with probability p, and that the redundancy is linearly independent. Then, in the case of failure, one can run UDS on the redundancy, and if this succeeds (with probability $p' > 0$, due to linear independence), the target state can still be generated from the outcome. This overall procedure ($\mathcal{T}$ followed by UDS in case of failure) comprises a uniform probabilistic transform $\mathcal{T}'$ which performs the same task as $\mathcal{T}$ but succeeds with probability $p' + p > p$. Hence $\mathcal{T}$ could not have been optimal.

10.1.3 Transformations between symmetric sets of pure states

As noted, the case when the sets of states in focus is symmetric is of interest since many quantum protocols [86, 136, 1, 3, 137]) work with symmetric quantum states.

A set of (pure) states $A = \{|a_i\rangle\}_{i=0}^{N-1}$ is *symmetric* if there exists a fixed unitary U with the property

$$U|a_i\rangle = |a_{(i+1) \bmod N}\rangle \text{ for all } i. \tag{10.4}$$

The symmetry addressed above and in the remainder of this section is a cyclic symmetry. The assumption that source and target states are symmetric allows us to link probabilistic and uniform probabilistic transforms. In this case, any probabilistic transform can be 'uniformized', as shown by the following lemma:

Lemma 31. *(Uniformization) If there exists a probabilistic transformation taking the states in A to states in B, which succeeds with the probabilities $\{p_i\}_{i=1}^N$, where A and B are symmetric sets of states, then there exists a uniform probabilistic transform taking the states in A to states in B which succeeds with probability*

$$p = \frac{1}{N} \sum_{i=1}^{N} p_i. \tag{10.5}$$

The proof of this lemma is given in Section 10.3.

Additional properties of uniform transforms with symmetric source and target states are rooted in the structural properties of Gram matrices of sets of symmetric states:

[2] This result was proven using different techniques and stated in a different formalism in [131].

Lemma 32. *A Gram matrix of kets is a circulant matrix if and only if the corresponding set of kets is symmetric.*

Proof of this lemma is given in Section 10.3.

A circulant matrix is a square matrix, defined by its first row, for which the i^{th} row is the right-circular shift of the first row by $i-1$ positions. Circulant matrices frequently appear in signal processing, and have two convenient properties: $i)$ circulant matrices diagonalize when conjugated by the unitary discrete Fourier transform (DFT) matrix, and $ii)$ the discrete Fourier transform of the first row of the circulant matrix is a vector containing the eigenvalues of the circulant matrix [138]. The discrete Fourier transform matrix of size N is the Vandermonde matrix of the N^{th} primitive roots of unity, given with

$$DFT = \left[\exp \frac{-2(p-1)(q-1)i\pi}{N}\right]_{p,q}, \;\; p = 1, \ldots N, q = 1, \ldots N \tag{10.6}$$

which, when scaled by the pre-factor $1/\sqrt{N}$ becomes unitary, and which we then denote $uDFT$.

The criterion for the existence of a uniform probabilistic transform taking states from A to B, succeeding with probability p, is the existence of Gram matrices of states Π^s and Π^f such that the equation

$$G_A = p\Pi^s \circ G_B + (1-p)\Pi^f \tag{10.7}$$

holds. This is a slight simplification of the more general condition in Lemma 30.

In general, probabilistic uniform transforms with symmetric source and target sets may have leak and redundancy which are not symmetric. Nonetheless, the following lemma shows that such a transform has a variant with the same success probability where the leak and redundancy are symmetric:

Lemma 33. *(Symmetrization) If there exists a uniform probabilistic transform taking states from a symmetric set A to a set of symmetric states B, succeeding with some probability p, then there exists a uniform probabilistic transform taking the states from A to B, succeeding with probability p, where the leak and redundancy are symmetric.*

Proof of this lemma is given in the the Section 10.3.

10.1.3.1 Finding optimal uniform transforms

Both from a practical and theoretical point of view, when considering transforms from a source to a target set one is often most interested in the optimal transforms. Optimality is naturally defined only in the case of uniform transforms. However, by virtue of Lemma 31, when transforms with symmetric source and target sets are concerned, if any kind of transform linking the source and target states exists, then so does a uniform transform. In this sense, for transforms between symmetric sets, optimality can in principle always be defined as the optimality of the uniform

transform [3].

In general, given two sets of states A and B, the quest for the optimal uniform transform taking the states in A to states in B reduces to the maximization of the success probability p over the space of all positive-semidefinite matrices (of the appropriate size) Π^s and Π^f with unit diagonal, subject to the constraint given in expression (10.7). The dimensionality of the search space is then quadratic in the number of states. However, if source and target states are symmetric, as a consequence of Lemma 33, we may assume that Π^s and Π^f are circulant as well. Then all the matrices appearing in expression (10.7) are circulant, as the Hadamard product of circulant matrices is also circulant. Hence, they all diagonalize in the same basis, and the dimensionality of the search space reduces quadratically from $O(N^2)$ to $O(N)$, where N is the number of states.

The problem of finding optimal uniform transforms which have symmetric source an target sets is resolved by the following canonical linear program:

$$\begin{aligned} \text{maximize} \quad & \overrightarrow{c}^{\mathrm{T}}.\overrightarrow{x} \\ \text{subject to} \quad & M.\overrightarrow{x} \leq \overrightarrow{b} \\ \text{and} \quad & \overrightarrow{x} \geq 0, \end{aligned}$$

where $\overrightarrow{c}^{\mathrm{T}} = [1, \ldots, 1]$, $\overrightarrow{b} = \lambda_{G_A}$, and $M = DCM_{\lambda_{G_B}}$, which is a circulant matrix, where the i^{th} *column* is the vector λ_{G_B} 'downward' shifted by $i-1$ positions (the *discrete convolution matrix* $DCM_{\lambda_{G_B}}$ of the vector λ_{G_B}). The optimal success probability is given by

$$p = \frac{\overrightarrow{c}^{\mathrm{T}}.\overrightarrow{x}}{N}, \tag{10.5}$$

where the dot '.' (e.g. $\overrightarrow{x}^T.\overrightarrow{y}$ or $M.\overrightarrow{x}$) denotes the standard matrix product. The vector of eigenvalues of the Gram matrix of the leak of the optimal transform is given by $\lambda_{\Pi^s} = \frac{1}{p}\overrightarrow{x}$. As both G_A and G_B are circulant matrices, the vectors of eigenvalues λ_{G_A} and λ_{G_B} are computed by taking the discrete Fourier transform of the first row of G_A and G_B, respectively.

In the remainder of this section we show that the linear program above solves the problem of finding optimal uniform transforms. The constraint (10.7) where all the matrices are circulant can be written in terms of the vectors of eigenvalues of the matrices appearing, as they all diagonalize in the same basis:

$$\lambda_{G_A} = p\lambda_{\Pi^s \circ G_B} + (1-p)\lambda_{\Pi^f}. \tag{10.6}$$

Note that for the vector λ_{Π^f} to be a vector of eigenvalues of a circulant Gram matrix of states, it is sufficient and necessary that all its entries are non-negative and sum up to N. Using the

[3] One may be tempted to do the same for non-symmetric transforms. However there exist non-uniform transforms which have non-symmetric source and/or target sets which succeed with non-zero probability for some states at least, for which no uniform transform exists (all uniform transforms fail with unit probability).

circular convolution Theorem it can be shown that

$$\lambda_{\Pi^s \circ G_B} = \lambda_{\Pi^s} * \lambda_{G_B}, \tag{10.7}$$

where $*$ represents the (normalized) discrete convolution (or discrete cross-correlation) of vectors defined al follows. If $\overrightarrow{x}$ and $\overrightarrow{y}$ are two vectors of size N, with corresponding entries x_i and y_i for $i = 0, \ldots, N-1$, then $\overrightarrow{z} = \overrightarrow{x} * \overrightarrow{y}$ is a length N (with components denoted z_i), defined component-wise by

$$z_i = \frac{1}{N} \sum_{j=0}^{N-1} x_j y_{[(N-j+i) \bmod N]}. \tag{10.8}$$

The discrete convolution of two vectors can also be represented in terms of a matrix-vector product by using the discrete convolution matrix $DCM_{\overrightarrow{x}}$ of the vector $\overrightarrow{x}$ defined via its transpose: the transpose matrix $DCM_{\overrightarrow{x}}^T$ is a circulant matrix whose first row is the vector $\overrightarrow{x}$. It holds that $\overrightarrow{x} * \overrightarrow{y} = DCM_{\overrightarrow{x}}.\overrightarrow{y} = DCM_{\overrightarrow{y}}.\overrightarrow{x} = \overrightarrow{x} * \overrightarrow{y}$. Hence, the constraint (10.6) is equivalent to

$$\lambda_{G_A} = p\, DCM_{\lambda_{G_B}} \lambda_{\Pi^s} + (1-p)\, \lambda_{\Pi^f} \quad , \tag{10.9}$$

which can be shown to be equivalent to the inequality

$$\lambda_{G_A} - p\, DCM_{\lambda_{G_B}} \lambda_{\Pi^s} \geq 0, \tag{10.10}$$

where λ_{Π^s} is a non-negative real vector, whose entries sum up to N. The inequality above is interpreted component-wise [4]. To obtain the linear program stated at the beginning of this section, we note that if a vector $\overrightarrow{x}$ is a vector of length N with non-negative entries $\{x_i\}$, which maximizes $s = \sum_{i=1}^{N} x_i$ subject to the constraint

$$DCM_{\lambda_{G_B}} \overrightarrow{x} \leq \quad \lambda_{G_A}, \tag{10.11}$$

then $\lambda_{\Pi^s} = \dfrac{N}{s} \overrightarrow{x}$ allows for the maximal p subject to constraint (10.10), and the maximum is reached at $p = \dfrac{s}{N}$.

10.1.4 The geometric interpretation of the optimization procedure

As we have shown, the search for the optimal probability of success p_{opt} of a uniform transform which takes N input states to N output states, where both sets of states are symmetric, reduces

[4] To prove this equivalence, note that (10.9) implies the constraint (10.10) as the eigenvalues of Π^f have to be non-negative. To see that the inverse holds as well, it suffices to show that if λ_{Π^s} is a vector of positive components which sum up to N, then the entries of the vector $\lambda_{G_A} - p\, DCM_{\lambda_{G_B}} \lambda_{\Pi^s}$ sum up to $(1-p)N$. By construction, the entries of λ_{G_A} sum up to N. Recall that $DCM_{\lambda_{G_B}} \lambda_{\Pi^s}$ is also a vector of eigenvalues of a Gram matrix of a symmetric set of kets. Hence its components also sum up to N. Hence, the components of $\lambda_{G_A} - p\, DCM_{\lambda_{G_B}} \lambda_{\Pi^s}$ sum up to $N - pN = (1-p)N$, and we have shown the equivalence of constraints (10.6), (10.9) and (10.10).

to the following optimization problem:

p_{opt} is the maximal p subject to constraint

$$\lambda_{G_A} = p\, DCM_{\lambda_{G_B}} \lambda_{\Pi^s} + (1-p)\, \lambda_{\Pi^f} \quad . \tag{10.12}$$

where λ_{Π^s} and λ_{Π^f} are some non-negative real vectors, whose entries sum up to N.

We have also shown that the constraint above is equivalent to the inequality

$$\lambda_{G_A} - p\, DCM_{\lambda_{G_B}} \lambda_{\Pi^s} \geq 0, \tag{10.13}$$

where λ_{Π^s} is a non-negative real vectors, whose entries sum up to N.

The search space defined by the constraint (10.10) is the space of all points embedded in an N dimensional space whose coordinates sum up to N. This is a convex set, defined by the extreme points $\{e_i\}_{i=1}^N$, where e_i is a vector with the number N as the i^{th} component, and zeroes elsewhere. But then, by the linearity of matrix-vector multiplication, the set

$$S = \left\{ DCM_{\lambda_{G_B}} \lambda_{\Pi^s} | \lambda_{\Pi^s} \geq 0, ||\lambda_{\Pi^s}||_1 = N \right\} \tag{10.14}$$

is a convex set as well. The norm $||\cdot||_1$ is defined as the sum of the absolute values of the entries of the vector in the argument. It is easy to see that S is the convex hull of the columns of the matrix $N \times DCM_{\lambda_{G_B}}$. First, let us assume $DCM_{\lambda_{G_B}}$ is non-singular, which is equivalent to saying that the set of target kets does not contain mutually orthogonal kets. Then it holds that the columns of the matrix $N \times DCM_{\lambda_{G_B}}$ are also the extreme points of the set S. The constraint (10.10) can then be written as

$$\lambda_{G_A} - p\, X \geq 0 \tag{10.15}$$

where $X \in S$.

Let T be the set defined as follows:

$$T = conv(\{e_i\}_{i=1}^N), \tag{10.16}$$

where $conv(A)$ denotes the convex hull of the set of points A. T is the convex set of all points which correspond to a symmetric set of N states. This is a regular $(N-1)$-simplex which we can embed in the vector space $\mathbb{R}^N$. Clearly, the point λ_{G_A} is an element of T, and S is also a regular $N-1$-simplex, contained in T. Simplices T and S share their centre at coordinates $(1, \ldots, 1)$, and S is in a scaled down, rotated copy of T. From this it can be shown that the rows of the matrix $DCM_{\lambda_{G_B}}$ are the extreme points of the set S even when the discrete convolution matrix is singular. The only exception is the degenerate case when all the entries of $DCM_{\lambda_{G_B}}$ are equal, which corresponds to the case when the set of target states is an orthogonal basis.

It is easy to see that, if λ_{G_A} lies in S, then the constraint (10.10) can be satisfied for $p = 1$, *i.e.* there exists a deterministic transform from the set of states A to the set of states B. If this is not

the case then the geometric interpretation of the constraint is as follows:

Lemma 34. *For a $0 < p \leq 1$ there exists a solution X satisfying the constraint (10.15) if the intersection between the simplex $p \times S = \{p \times \overrightarrow{x} | \overrightarrow{x} \in S\}$, embedded in $\mathbb{R}^N$, and the (N) orthotope (hyperrectangle or box) L defined by the opposite points $(0, \ldots, 0)$ and λ_{G_A} is not the origin point alone.*

The orthotope L can be defined as

$$L = \left\{ \overrightarrow{x} \in \mathbb{R}^N | \lambda_{G_A} - \overrightarrow{x} \geq 0 \right\}, \tag{10.17}$$

which makes the validity of the geometric interpretation above obvious. The geometric interpretation is illustrated in Figure 10.1, for the case $N = 3$.

Let us now consider a few special cases, for illustration purposes. It is clear that if $S = T$ then for any set of input states the transform can be done deterministically, as $\lambda_{G_A} \in T$. However, if $S = T$, then $DCM_{\lambda_{G_B}}$ has exactly one '1' in each row and each column, hence the vector of eigenvalues λ_{G_B} has one entry equal to N and the rest is zero. This corresponds to the setting where the target set of states comprises exactly one state, and the deterministic transformation preforming this is the contraction to that particular state.

In the opposite scenario, S may consist of a single point – the point $(1, \ldots, 1)$. In this case $DCM_{\lambda_{G_B}}$ is a matrix containing just unities, and the corresponding set of target states is then orthogonal. For there to exist a solution satisfying the constraint (10.15), by the geometric interpretation, the line $\{p(1, \ldots, 1) | 0 < p \leq 1\}$ must intersect the orthotop L. This happens if and only if the extreme point λ_{G_A} which defines the orthotope has all components non-zero. This requirement implies that the input set of states is linearly independent. If we recall that a transformation with orthogonal target states is equivalent to unambiguous discrimination of input states, then we see we have recaptured a well-known result [5]: a set of states can be unambiguously discriminated if and only if the set is linearly independent.

Finally we can use the geometric interpretation to give a new result, which we haven't addressed thus far: If a uniform transformation with symmetric sets of input and output states is optimal, then the leak is linearly dependent. To show this we will adopt a dynamic picture as illustrated in Figure 10.2. Let $\lambda_{G_A} \notin T$. What we seek is the largest p such that the simplex $p \times S = \{p \times \overrightarrow{x} | \overrightarrow{x} \in S\}$ and the orthotope L intersect. The simplex $p \times S$ clearly lies in the simplex $p \times T = \{p \times \overrightarrow{x} | \overrightarrow{x} \in T\}$, and the intersection will occur in this simplex. As we slowly decrease p the intersection between the simplex $p \times T$ and the orthotope L 'grows' while the simplex $p \times S$ slowly reduces in size. At one point, for some p, the intersection of the simplex $p \times T$ and the orthotope L touches the simplex $p \times S$, if there is a solution to the problem. Whenever this happens, the touching point is clearly on the face of the simplex $p \times S$, and not an interior point.

This means that the solution (the corresponding touching point in S) is a convex combination of at most $N - 1$ rows of $DCM_{\lambda_{G_B}}$, which in turn implies that λ_{Π^s} has a zero component. Since

[5] Restricted, however, to symmetric sets of input states.

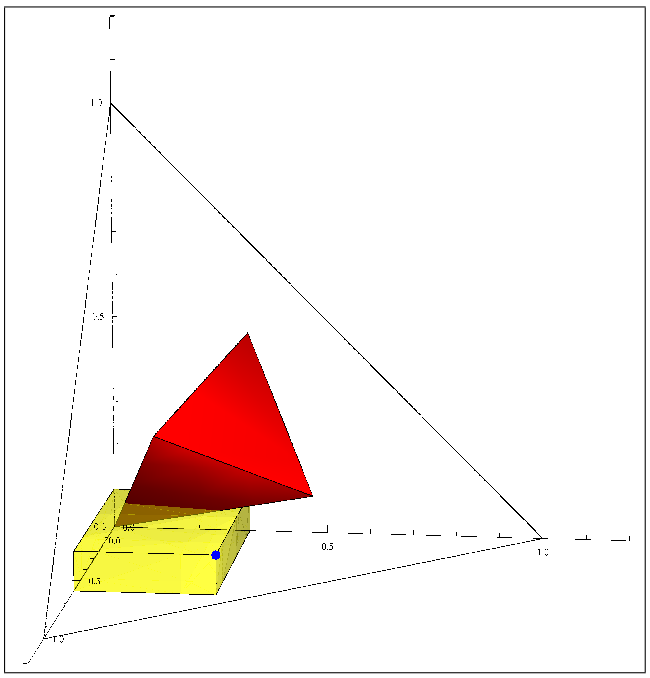

Figure 10.1. Illustration of the geometric interpretation of the solution existence criterion given in expression (10.15), for $N = 3$. The the vector of eigenvalues of the Gram matrix of the source states is represented as a single point (blue in our illustration) which lies somewhere in the simplex defined by the extremal points $(0, 0, 1), (0, 1, 0), (1, 0, 0)$. This simplex is represented by the transparent triangle. The source states (represented by the blue point) uniquely define the orthotop L – a box, given in yellow. All points in the orthotope (and only those points) have the property that $i)$ all their components are non-negative, and $ii)$ any point in the orthotope when subtracted coordinate-wise from the blue point (defined by the source states) gives a point with non-negative components. The vector of eigenvalues of the Gram matrix of the target states defines the corresponding discrete convolution matrix, the columns of which are the extremal points of the search space S defined at the beginning of this section. The space S is, for $N = 3$, a regular 2-simplex (moreover, an equilateral triangle), embedded in a 3-dimensional space, and is represented with the red triangle which lies within the transparent triangle representing all possible vectors of eigenvalues of the Gram matrix of a size three symmetric set of states. Together with the origin, the extremal points of S define the 3-simplex $p \times S = \{p \times x | x \in S\}$ which is represented by the entire red tetrahedron in the illustration. Lemma 34 states that the necessary and sufficient criterion for the existence of a probabilistic transform taking the source to the target states is that the intersection between the red tetrahedron and the yellow box is not just the point of origin. The point of origin would correspond to a transform which succeeds with probability zero. Note also that clearly the intersection of the orthotope and the tetrahedron can be just the point of origin only if the orthotope has at least one dimension zero. This corresponds to the setting where source states are linearly dependent.

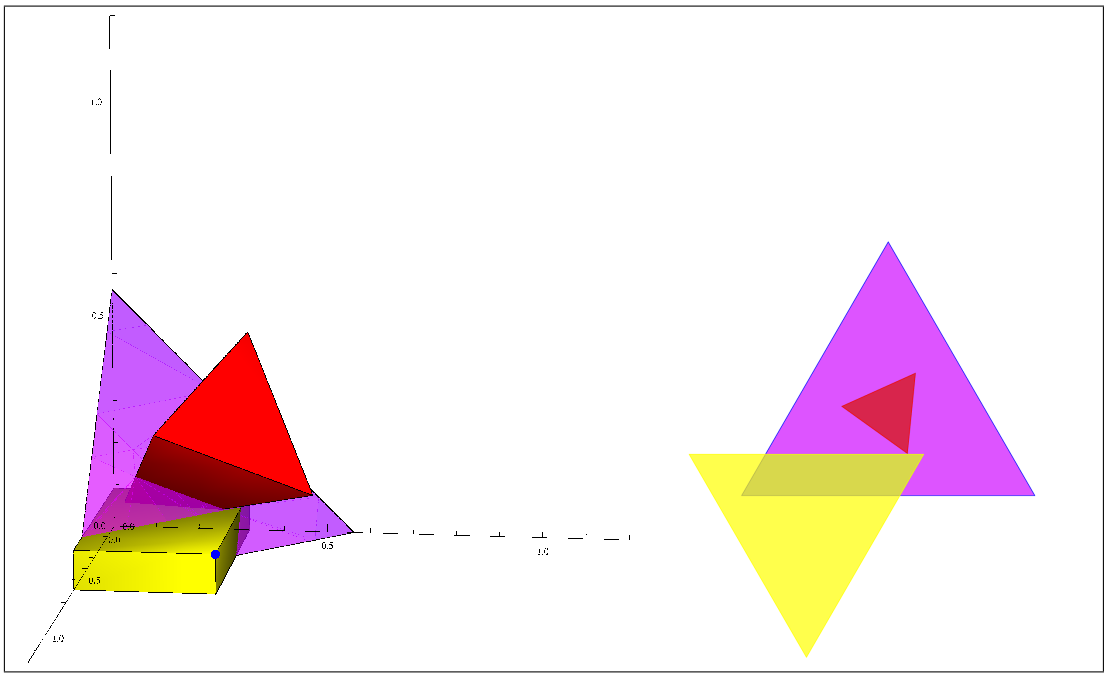

Figure 10.2. Dynamic picture: The left-hand side illustration represents the setting where the optimal transform has been found. For a particular (maximal) p *the simplex which is the intersection of the (red) tetrahedron and the simplex* $p \times T = conv((0,0,p),(0,p,0),(p,0,0))$ *represented with the purple triangle touches the intersection of the yellow orthotope* L *and the same simplex* $p \times T$*. The right-hand side image illustrates this event when the view is restricted only to the* $N-1=2$*-dimensional unique hyperplane* H_p *containing the simplex* $p \times T$*. The intersection of the hyperplane* H_p *and the whole red tetrahedron makes up the red triangle, the purple triangle is the convex set* $p \times T = conv((0,0,p),(0,p,0),(p,0,0))$ *in the same hyperplane. The yellow triangle corresponds to the intersection of the hyperplane* H_p *and the extended orthotope* L' *where the edges of the original orthotope are allowed to extend to* $-\infty$ *each. This intersection is a regular 2-simplex again, with a fixed orientation with respect to the simplex* $p \times T$*. In the dynamic picture, if we were to let* p *slowly decrease from unity, we would witness the yellow triangle emerge from a single point, and grow until it touches the red triangle. The centre of the yellow triangle would slowly move towards the closest extremal point of* $p \times T$ *due to the change of its position in the barycentric coordinate system of the simplex* $p \times T$ *which changes as a function of* p*.*

λ_{Π^s} is the vector of eigenvalues of the Gram matrix of the leak, this means the leak is linearly dependent. If we now join this with the fact that optimal transforms have a linearly dependant redundancy, shown in Section 10.1.2.1, we get the following statement:

Lemma 35. *If a uniform transformation with symmetric sets of input and output states is optimal, then the leak and the redundancy are linearly dependent.*

The inverse however, does not hold.

10.1.5 Geometric characterisation of the leak and the redundancy

As we have shown, a uniform transform from symmetric to symmetric states succeeding with the probability p can always be realized in such a way that the leak and the redundancy are symmetric sets of states (Lemma 33). In this case, the leak and the redundancy can completely be characterised from the geometric picture. Recall that, if the transform exists for a fixed p, then the intersection between the simplex $p \times S$ and the orthotop L is non-empty and this intersection is contained in the simplex $p \times T$. It is easy to see that the intersection $p \times F = L \cap p \times T$ is a convex set, more precisely, a bounded convex polytope.

Let X be a solution, obeying the constraint (10.15). The vector X completely characterises the leak. Recall, the vector X is of the form $DCM_{\lambda_{G_B}}\lambda_{\Pi^s}$, where λ_{Π^s} is the vector of eigenvalues of the leak set. Thus, the vector X, viewed as a point in the simplex $p \times T$ embedded in the Euclidean space $\mathbb{R}^N$, is a convex combination of the rows of the matrix $DCM_{\lambda_{G_B}}$. The weights of this convex combination are the components of λ_{Π^s}. In other words, the representation of X in the barycentric coordinates given by the extreme points of $p \times S$ (these points are the rows of the (scaled) matrix $DCM_{\lambda_{G_B}}$) gives exactly the vector of eigenvalues of the Gram matrix Π^s. A barycentric coordinate system is a coordinate system in which a point's position is specified as the centre of mass, or barycenter, of masses placed at the vertices of a simplex, in our case the simplex $p \times S$, which is the convex hull of the rows of the matrix $p \times DCM_{\lambda_{G_B}}$.

An analogous observation can be done for the redundancy – X represented in the barycentric coordinates of some of the the extreme points of $L \cap p \times T$ [6] will give us the structure of the redundancy. While this relationship is more involved than in the case of the leak, and we leave it for further research, certain easy observations can be made for the optimal transform case.

As we noted, if the transform is optimal, then the solution point X lies in the intersection of the faces of the polytope $p \times F = L \cap p \times T$ and the simplex $p \times S$, *i.e.* it is not in the interior of either. If the dimensionality of the face which is involved in the contact of $p \times S$ is zero (a vertex) then every symmetrized optimal transform is always leakless. Similarly, if the dimensionality of the face which is involved in the contact of $p \times F$ is zero, then it is redundancy-free.

If the contact involves faces of higher dimensionalities of $p \times F$ then essentially anything may happen, depending on the structure of the overlap. In the example given in the right-hand side illustration of Figure 10.2, the contact point for the simplex $p \times S$ (red) and the simplex $p \times L' \cap H_p$

[6] The number of the extreme point of this polytope may be larger than $N+1$, but by Carathéodory's theorem, each point in this polytope can be represented as a convex combination of at most $N+1$ points.

(yellow) is a vertex of the simplex $p \times S$, and thus this transform is leakless. However, the contact point is interior of a 1-dimensional face of the yellow simplex, indicating that the vector of eigenvalues of the redundancy has two non-zero entries. Thus, the redundancy comprises at least two non-equal vectors.

Note that the structure of the overlap depends on the relative orientations and positions of the polytope $p \times F$ and the simplex $p \times S$. As we noted, the simplex $p \times S$ is just a scaled down and rotated simplex $p \times T$. The orientation of the polytope $p \times F$ is in a sense fixed with respect to the orientation of $p \times T$. To explain this, consider the simplex $L' \cap p \times T'$ where $L' = \prod_{i=1}^{N} \left\langle -\infty, \lambda^i_{G_A}\right]$, $\lambda^i_{G_A}$ being the i^{th} component of the vector λ_{G_A} and the product is the Cartesian product. We define pT' to be the hyperplane defined by the points $\{p \times e_i\}_{i=1}^{N}$. The set L' is just the extended orthotope L where the sides (1-faces) radiating from the point λ_{G_A} are allowed to stretch to $-\infty$. Then $L\prime \cap p \times T'$ is the intersection of $L' \cap p \times T'$ and the positive quadrant $\prod_{i=1}^{N} [0, \infty\rangle$. $L' \cap p \times T'$ is then a regular N-simplex, and if we translate it by moving the centre to the point $(p, \ldots, p)$ we have a simplex which is a scaled, centrally mirrored copy of $p \times T$. In this sense, the orientation of $p \times F$ (recall, $p \times F = (L' \cap p \times T') \cap \prod_{i=1}^{N} [0, \infty\rangle$) is fixed, relative to the orientation of $p \times T$.

10.1.5.1 Quantifying the leak and the redundancy

Transformations between different types of quantum states become unavoidable when heterogeneous encodings are used for different aspects of quantum information tasks. In particular, such transform may be part of a cryptographic protocol, in which case quantifying the leak and redundancy in information-theoretic terms becomes crucial. For instance, one can imagine a simple two-party scheme in which party A, traditionally called Alice, wishes to send to party B, called Bob, information encoded in quantum states comprising the set of target states B. However, Alice has at her disposal only quantum states from a set of quantum states A. So, Alice indeed does send her information encoded as states in A to Bob, who performs an optimal probabilistic transform in order to obtain the target state B. For example, an ideal protocol may call for single-qubit states, but Alice can only generate pure states which approximate qubit states. It is then important for Alice to know what additional information Bob can obtain when transforming source states to target states [7]. As we have seen, such a transform is characterised by an expression of the form $G_A = p\Pi^s \circ G_B + (1-p)\Pi^f$ where the Gram matrices G_A and G_B fully characterise the source and target states (up to unitary equivalence), and Π^s and Π^f characterise the *leak* and the *redundancy*, that is, the residual states when the transform succeeds and when it fails, respectively. One way by which Alice may quantify the leak of information (embodied in the leak states) is by calculating the accessible information in this set of states. If $\{\rho_i\}_{i=0}^{N-1}$ is a set of quantum states, then the accessible information I_{acc} in this set of states is bounded above by the Holevo χ quantity, $I_{acc} \leq \chi(\rho_{AVG}) = S(\rho_{AVG})$, where $\rho_{AVG} = 1/N \sum_i \rho_i$ is the average state if each ρ_i appears equally likely as a message, $S(\cdot)$ denotes the Von Neumann entropy,and

[7] Such approximations often appear in many proposals for realizations of quantum cryptographic protocols: polarization-encoded photons (which realize a qubit) are often approximated by polarized weak coherent pulses. In this case, almost without exception, a new security analysis is required.

the last equality holds if ρ_i are pure. If $\lambda = (\lambda_0, \ldots, \lambda_{N-1})$ is the vector of the eigenvalues of ρ_{AVG}, then the Von Neumann entropy can be expressed in terms of the Shannon entropy H as $S(\rho_{AVG}) = -\sum_{i=0}^{N-1} \lambda_i \log \lambda_i$.

If $A = \{|a_i\rangle\}_{i=1}^{N}$ is a set of kets (pure states), then using matrix algebra it can be shown that the non-zero eigenvalues of the matrix G_A and the operator $\sum_i |a_i\rangle\langle a_i|$ are equal. Hence, the upper bound on the accessible information in a set of states can be calculated as the Shannon entropy of normalized eigenvalues of the Gram matrix of that set of states. The optimization procedure we have presented, which finds the optimal success probability p, also finds the corresponding vector λ_{Π^s}. From this, λ_{Π^f} is easily computed, which are the eigenvalues of the Gram matrices of the leak and of the redundancy. From these eigenvalues it is then very simple to directly upper bound the accessible information in the leak and the redundancy.

10.1.6 Application: From coherent states to qubit states

Traditionally, for most applications of quantum information processing, the information is encoded in qubit states. However, it is also possible to use continuous-variable states, that is, states of the quantum harmonic oscillator (e.g. coherent states). In this section the source states will be a set of coherent states

$$A = \left\{|a_k\rangle = |e^{i\theta_k}\alpha\rangle\right\}_{k=0,\ldots,N-1} \tag{10.18}$$

where α is a real amplitude and θ_k are their phases. The target states are the qubit states in the Bloch sphere XY plane,

$$B = \left\{|b_k\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + e^{i\theta_k}|1\rangle\right)\right\}_{k=0,\ldots,N-1}. \tag{10.19}$$

By choosing the angles θ_k as $\theta_k = 2k\pi/N$ we obtain a very common family of encodings, which incidentally renders the sets A and B symmetric.

The problem we resolve is finding the optimal uniform transform taking the states in the set A to those in B. Initially, let us assume N is even. We may immediately note that the states in A are linearly independent, so an unambiguous measure-and-prepare process will get us the desired transform succeeding with the success probability of an UDS procedure applied on the states in A. The optimal success probability of such a UDS procedure establishes a lower bound, and an upper bound is found by noting that if N is even, then the desired probabilistic transform maps any two input states with relative phases differing by π into orthogonal states. Hence, in particular this transform effectively performs unambiguous discrimination of the states $|\alpha\rangle$ and $|-\alpha\rangle$. By using the results of Section 10.1.2.1, the success probability of this UDS procedure (hence of the overall probabilistic transform) is upper bounded by $s_{bound} = 1 - \exp(-2\alpha^2)$. This bound is always higher than the probability of unambiguous discrimination, except for the case of two states, where they coincide. The cases for 4 and 8 states are illustrated in Figure 10.3.

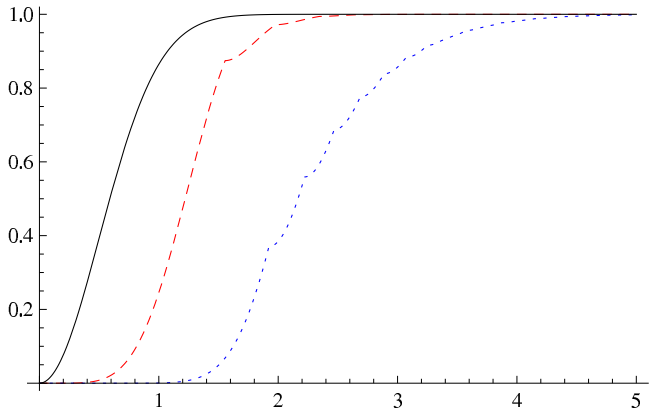

Figure 10.3. Comparison of the optimal success probability of unambiguous discrimination of 4 (red, dashed) and 8 (blue, dotted) states of a symmetric set of states, as a function of the real amplitude α. The black curve represents the optimal success probability of the coherent to qubit states transform, which is independent of the number of states.

In the remainder of this section we prove, constructively, that the upper bound can always be reached. This is done by first obtaining results for the case $\alpha \leq 1$, and then using these results, constructing transforms also for the case $\alpha > 1$.

To begin, we introduce the notion of a multiprobabilistic transform, defined in [133]. Multiprobablistic transforms are a generalization of probabilistic transforms, where there may be many different sets of targets states and with some probabilities an input state is transformed to a corresponding state in one of the target sets. For our purposes, we shall define the uniform version of such transforms:

Definition 36. *Let $S = \{|s_i\rangle\}_{i=1}^n$ be a set of source states and $T^j = \left\{|t_i^j\rangle\right\}_{i=0}^n$ for $j = 0, \ldots, k-1$ be a collection of possible target states. A uniform multiprobabilistic transform $\mathcal{T}$ from the set S to the sets in $\{T^j\}_j$, succeeding with the probability vector $(p_0, \ldots, p_{k-1})$, where $\sum_{i=0}^{k-1} p_i = 1$ and for all i $p_i \geq 0$, performs*

$$\mathcal{T}(|s_i\rangle) = |t_i^j\rangle \text{ with probability } p_j \tag{10.20}$$

for $i = 1, \ldots, n$ and $j = 0, \ldots, k-1$.

The set T^0 corresponding to success probability p_0 is reserved for the 'fail outcome' states, analogous to the redundancy set of states in probabilistic transforms.

As a consequence of Theorem 3 in [133], for fixed source set S and target sets $\{T^j\}_{j=1}^{k-1}$ and a probability vector $(p_0, \ldots, p_{k-1})$, such a uniform transform exists if and only if there exists a set of Gram matrices of states $\{\Pi^f, \Pi^1, \ldots, \Pi^{k-1}\}$ such that the following equality holds:

$$G_S = p_0\Pi^f + p_1 G_{T^1} \circ \Pi^1 + \cdots + p_{k-1} G_{T^{k-1}} \circ \Pi^{k-1}, \tag{10.21}$$

where G_S is the Gram matrix of the set S and G_{T^j} the Gram matrix of the set T^j for all j. We will call such a transform leakless if the matrices $\Pi^j = \mathbf{1}$ for all j are matrices with all entries

being the unity, and redundancy-free if the matrix $\Pi^f = \mathbf{1}$ [8].

Let us now define a collection of sets of target states B^j as

$$B^j = \{|b_i\rangle^{\otimes j}\}_{i=0}^{N-1}, j = 1, 2, \ldots, N-1. \tag{10.22}$$

That is, the set B^j comprises states which are j-fold copies of the elements of the (original target) set $B \equiv B^1$, which are the XY plane qubit states. Then we have the following lemma, which holds specifically for the source and target states of interest:

Lemma 37. *Let the amplitude α of the states in the set A, defined in equation (10.29), satisfy $0 < \alpha \leq 1$. Then there exists a uniform multiprobabilistic transform with the success probability vector $(p_0, \ldots, p_{N-1})$, which takes the states from the set A to the collection of target states $\{B^j\}_{j=1}^{N-1}$* and *is redundancy-free and leakless. The failure probability p_0 of this transform is equal to $\exp(-2\alpha^2)$.*

The proof of this Lemma is somewhat cumbersome and is presented in Section 10.3. The requirement that the transform be redundancy-free and leakless uniquely fixes the transform, up to the freedom in the choice of the realized fixed failure-outcome states.

As a corollary of this Lemma we obtain the desired uniform probabilistic transform from coherent states in A, for $0 < \alpha \leq 1$, to the qubit states in B, and a characterisation of the leak and redundancy for this optimal transform, as we now show.

Corollary 2. *Let A and B be symmetric sets of an even number N states, as defined at the beginning of this section, and let $0 < \alpha \leq 1$. Then there exists a redundancy-free uniform probabilistic transform taking the states from A to corresponding states in B succeeding with probability $p_{succ} = 1 - \exp(-2\alpha^2)$. This transform is also optimal.*

Proof:

Lemma 37 establishes the existence of a multiprobabilistic uniform transform from the set A to the sets $\{B^j\}_{j=1}^{N-1}$, which is both redundancy-free and leakless, when $0 < \alpha \leq 1$. But then, by Theorem 3 in [133] there exists a probability vector $(p_0, \ldots, p_{N-1})$ such that the following equality holds:

$$G_A = p_0\mathbf{1} + p_1 G_{B^1} + \cdots + p_{N-1} G_{B^{N-1}}. \tag{10.23}$$

Note, the expression above is the necessary and sufficient condition given in expression (10.21) for the existence of a uniform multiprobabilistic transform, which is now both redundancy-free, and leakless.

Since the Hadamard product is distributive, and by expression (10.54), this expression (10.23) can be rewritten as

$$G_A = p_0\mathbf{1} + (1-p_0) G_{B^1} \circ \left(\frac{p_1}{1-p_0} G_{B^0} + \cdots + \frac{p_{N-1}}{1-p_0} G_{B^{N-2}} \right), \tag{10.24}$$

[8] Note that $\mathbf{1}$ is a Gram matrix of any set of unit vectors which are all equal.

with $G_{B^0} = \mathbf{1}$. Let us denote expression in the parenthesis in the equation above by Π^s,

$$\Pi^s = \frac{p_1}{1-p_0}\mathbf{1} + \frac{p_2}{1-p_0}G_{B^1} + \cdots + \frac{p_{N-1}}{1-p_0}G_{B^{N-2}}. \tag{10.25}$$

Note that Π^s is a Gram matrix of states, as it is a convex combination of Gram matrices of states. So we have

$$G_A = p_0\mathbf{1} + (1-p_0)G_{B^1} \circ \Pi^s. \tag{10.26}$$

This expression is a sufficient criterion for the existence of a uniform probabilistic transform taking the defined coherent states to qubit states. Since the fail probability is $p_0 = \exp(-2\alpha^2)$, by the upper bound on the success probability derived at the beginning of this section, it is the lowest possible, and this transform is optimal. This transform is also redundancy-free, as the Gram matrix of the redundancy is $\mathbf{1}$, that is, a Gram matrix of a set comprising identical states. The leak of this transform is symmetric by Lemma 32, as the matrix Π^s is a weighted sum of circulant matrices (see expression (10.25)), hence circulant itself. □

By investigating the expression (10.25), we can construct the leak states of this transform explicitly. The leak state $|\psi_i\rangle$, corresponding to the input state $|a_i\rangle$ can, up to unitary equivalence, be written as

$$|\psi_i\rangle = \sum_{j=0}^{N-2} \sqrt{\frac{p_{j+1}}{1-p_0}} |b_i\rangle^{\otimes j} \otimes |0\rangle^{\otimes N-2-j} \otimes |j\rangle \tag{10.27}$$

where the states of the last register (the *indicator* register) are orthogonal for differing labels, and we define for any state $|\eta\rangle$, the zeroth tensoral power $|\eta\rangle^{\otimes 0} \equiv 1$ (the unity of the field underlying the Hilbert space, i.e. the number one). These 'leaky' states are superpositions of varying numbers of copies (from zero to $N-2$) of the target state $|b_i\rangle$, all living in orthogonal subspaces of a larger Hilbert space (due to the orthogonality of the indicator register states).

We will now prove the existence of an optimal transform for any amplitude, also $\alpha > 1$. To do this we first note that coherent states can be 'split' into multimode states of a lower amplitude, i.e. there exists an isometry performing $U|e^{i\phi}\alpha\rangle = \bigotimes_{k=0}^{M-1} |e^{i\phi}\beta_k\rangle, \forall\, \phi$, as long as $\alpha^2 = \sum_{k=0}^{M-1} \beta_k^2$. We note that in quantum optics, this transform can be implemented by using balanced beamsplitters and phase shifters. Assume that we are given a set of coherent symmetric states A, as defined in equation (10.29) with $\theta_k = 2\pi k/N$, of amplitude $\alpha > 1$. Each of these states in A can be deterministically taken to the state $\bigotimes_{k=0}^{M-1} |e^{i\theta_k}\beta\rangle$ by 'splitting' the coherent state into M modes, where $\beta = \dfrac{\alpha}{\lfloor\alpha\rfloor + 1}$ and $M = (\lfloor\alpha\rfloor + 1)^2$. Now we have that $\beta \leq 1$ and $\alpha^2 = M\beta^2$, where M is a non-negative integer. By the Corollary 2 we have that each subsystem state $|e^{i\theta_k}\beta\rangle$ can be individually transformed to the corresponding qubit state in the set B with probability $\exp(-2\beta^2)$. Note that, if only one of the individual transforms performed on the states $|e^{i\theta_k}\beta\rangle$ succeeds, then we have succeeded in generating exactly one copy of the target state from the source state $|e^{i\theta_k}\alpha\rangle$. The probability of the transform failing on all M copies is $\exp(-2\beta^2)^M = \exp(-2\alpha^2)$. Hence, we have the following Theorem.

Theorem 9. *Let A and B be symmetric sets of an even number N states, as defined in equation (10.29) with $\theta_k = 2\pi k/N$, and let $\alpha > 0$. Then there exists a redundancy-free uniform probabilistic transform taking the states in A to the corresponding states in B, succeeding with probability $p_{succ} = 1 - \exp(-2\alpha^2)$. This transform is optimal.*

The leak of this overall transform will in general comprise multimode states, which in some modes contain a fixed state (the modes where the probabilistic transform failed), and in some modes the target qubit and the individual transform leak of the form given in expression (10.27). In contrast to unambiguous discrimination procedures for symmetric coherent states, the success probability of these optimal transforms generating qubit states does not depend on the number of states. In this analysis, we have assumed that the number of possible input states is even. As the success probability does not depend on the (even) number of states, the probabilistic transform can be done with the same success probability even when the number of states is N for an odd N. To see this, simply consider the transform which works for $2N$ states. The initial odd numbered symmetric states will be an interlaced subset of the extended set. However, here we do not have the validity of the upper bound any more, and it is not clear this success probability is optimal. While we do not offer a proof that the same bound holds for odd numbered states, evidence from performed numerical testing confirms this hypothesis.

An interesting aspect of the presented transform is that the success probability does not depend on the number of source and target states. Therefore it is possible that the same success probability may be reached when we consider the limit of an infinite number of states, $N \to \infty$. However, in the proofs of lemmas in this analysis, the finiteness of N is used, so proving this extension to the limit may be non-trivial. In the following section, we will however present a proposal for the realization of the presented transform, which does not assume a finite number of states, but achieves optimality in an asymptotic limit only.

10.1.6.1 Transforming coherent to qubit states using optical state truncation

After these results on the existence of optimal transforms, we will look at practical ways of implementing such transforms. A straightforward way of (sub-optimally) generating the desired qubit states from the source coherent states is through optical state truncation (OST) [139] or 'quantum scissors', as we will now describe. For a single mode state, such as a coherent state, OST is the probabilistic and heralded projection of the input state to a finite subspace (as defined by a selection of a number of Fock states), followed by renormalization of the state vector. OST has been realized using a linear optical network [140]. In this section we will focus on truncation to the subspace of the first two Fock states. Given the input state expanded in the number basis,

$$|\psi\rangle = \sum_{i=0}^{\infty} c_i |i\rangle,$$

where $\sum_i |c_i|^2 = 1$, OST is characterised by the POVM (POM) elements

$$\Pi_s = |0\rangle\langle 0| + |1\rangle\langle 1|, \;\; \Pi_f = I - \Pi_s \tag{10.27}$$

and, upon success, produces the state

$$|\psi_{trunc}\rangle = \mathcal{N}\left(c_0|0\rangle + c_1|1\rangle\right)$$

where the normalization factor is $\mathcal{N} = \left(|c_0|^2 + |c_1|^2\right)^{-1/2}$. If we now consider the input state to be a state from our source set of N coherent states,

$$|a_j\rangle := |e^{\theta_j i}\alpha\rangle = e^{-\alpha^2/2}\sum_{k=0}^{\infty}\frac{\alpha^k e^{k\theta_j i}}{\sqrt{k}}|k\rangle, \tag{10.27}$$

we see that the output state, after successful OST, which occurs with probability $p^{OST} = e^{-\alpha^2}(1+\alpha^2)$, is

$$|a_j{}^{OST}\rangle = \frac{1}{\sqrt{1+\alpha^2}}\left(|0\rangle + \alpha e^{i\theta_j}|1\rangle\right). \tag{10.28}$$

If $\alpha = 1$, this transform produces exactly the desired target qubit states.

This realisation does not, however, give the optimal success probability. The success probability of this transform for $\alpha = 1$ is approximately 0.735, which is less than the optimal value of approximately 0.864. The success probability of optical truncation to the vacuum and single photon subspace approaches unity more than exponentially quickly as the amplitude tends to zero. For $\alpha \neq 1$, the truncation will not produce the targeted qubit state, due to an uneven distribution of the weights between the $|0\rangle$ and $|1\rangle$ states. Re-weighting of the amplitudes can, however, also be achieved probabilistically, so now we consider the performance of the coherent to qubit transform realized by state truncation, followed by redistribution of the weights, for amplitudes $\alpha < 1$.

The redistribution of weights may optimally be done by applying a POVM defined by the positive elements

$$P_f = \gamma|0\rangle\langle 0|, \;\; P_s = I - P_f, \tag{10.29}$$

where $\gamma = 1 - \alpha^2$. These transforms fall into a class we call *umbrella transforms*. The success rate (the probability of outcome associated with P_s) of this transform is $p_{umb} = 2\alpha^2/(1+\alpha^2)$, hence the overall success probability of optical truncation followed by an umbrella transform for weight redistribution is

$$p_{overall} = p_{umb}p^{OST} = \frac{2\alpha^2}{1+\alpha^2}e^{-\alpha^2}(1+\alpha^2) = 2\alpha^2 e^{-\alpha^2}.$$

This value is always below the success probability of the optimal transform as the quotient

$p_{opt}/p_{overall}$ is equal to

$$\frac{p_{opt}}{p_{overall}} = \frac{\sinh(\alpha^2)}{\alpha^2}, \tag{10.29}$$

which is always greater than 1 on the interval of interest, approaching unity when $\alpha \to 0$.

10.1.6.2 Asymptotic optimality through beamsplitting

While the optimal transform of coherent to qubit states cannot be realized by OST followed by an umbrella transform to redistribute relative weights, it is evident that this transform performs better and better as the amplitude is reduced. It is natural to check whether a beamsplitting pre-procedure, analogous to the one used to prove the optimality Theorem 9 in the $\alpha > 1$, may be used to boost the overall success probability.

The procedure goes as follows: the input state of real amplitude α is 'beamsplitted' into M modes of amplitude $\alpha/\sqrt{M}$ with the same complex phase as the initial beam (as was done in the proof of Theorem 9). Then OST is applied to each of the beams, and if an individual OST succeeds, an umbrella transform is applied to re-weigh the vacuum and $|1\rangle$ components. The overall procedure succeeds if, for at least one of the split off beams, both the truncation and the umbrella transform are successful.

As we have shown, for a real amplitude α, a re-weighted OST produces the corresponding qubit state succeeds with probability $p_{overall} = 2\alpha^2 e^{-\alpha^2}$. Then, the success probability of the strategy where the input beam has been split into M beams is given by

$$p_{overall,M} = 1 - \left(1 - 2\frac{\alpha^2}{M} e^{-\alpha^2/M}\right)^M. \tag{10.30}$$

In the asymptotic case of infinitele many 'splits', the failure probability becomes

$$p_{overall,\infty} = \lim_{M\to\infty} \left(1 - 2\frac{\alpha^2}{M} e^{-\alpha^2/M}\right)^M = e^{-2\alpha^2},$$

which is equal to the failure probability of the optimal transform. The graph in Figure 10.4 compares the success probabilities of the optimal transform and the beamsplitter-assisted strategies for various numbers of splits M. This procedure can then arbitrarily well approach the optimal success probability. It is suitable for experimental realizations, as both quantum scissoring and the weight redistribution using umbrella transforms may be realized experimentally.

10.1.7 Conclusions

In this work we have addressed probabilistic transforms taking states from a 'source' to a 'target' set of quantum states, with emphasis on the case where these sets are symmetric. Such transforms can for example serve as interfaces between continuous-variable and finite-dimensional quantum

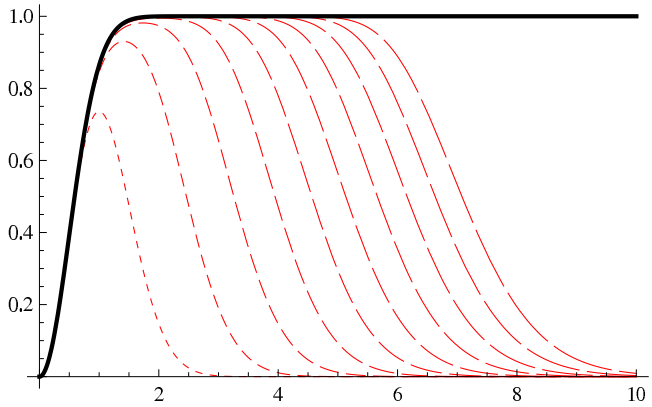

Figure 10.4. Comparison of the success probability of the optimal coherent to qubit states transform and the transform realized by beamsplitting into M beams of equal real amplitudes, followed by quantum scissors, followed by relative weight normalization between the vacuum and non-vacuum components on each of the weaker beams. The x axis gives the input amplitude α and y the success probabilities. The full (black) curve is the success probability of the optimal transform, and the (red) dashed curves the success probabilities of the beamsplitter-assisted quantum scissors strategies for $M = 1 \ldots 10$. The longer-dashed curves correspond to larger parameter M.

systems. State-dependent cloning and quantum state discrimination are also special cases of probabilistic transforms.

We have emphasised that in a probabilistic transform, information may be lost and leaked, which may have impact on the protocol efficiency or security. For this purpose we introduced the concepts of the leak and redundancy of a probabilistic transform. We have demonstrated how symmetric source and targets sets, which arise naturally in many quantum information applications, allows for a much simpler theory. In particular, we derived a linear program which finds optimal uniform probabilistic transforms in this symmetric setting. This constitutes a significant simplification over optimization techniques which must be employed in more general cases, and the dimensionality of the search space is reduced quadratically in the number of states considered. The presented method also allows for a simple characterisation of the aforementioned leak and redundancy.

Following this, we applied the derived theory to the problem of transforming a particular set of coherent states to a particular set of qubit states. Both sets appear in many quantum information protocols. The considered set of coherent states are so-called 'phase-locked' quantum states (e.g. used for quantum key distribution) suitable for long-range communication, and the set of qubit states is ubiquitous in quantum computation. For this setting, we derived the optimal transform and characterised the leak and the redundancy. By using beamsplitting, followed by the well-studied process of optical state truncation or 'quantum scissors', and an experimentally feasible amplitude re-weighing procedure, a probabilistic transform between these sets of states can be realized, albeit with sub-optimal success probability. The success probability of this procedure can however be made to asymptotically approach the optimal success probability.

An immediate application of such a transform may be in the realization of Universal Blind Quan-

tum Computation (UBQC) [1], in a case where the client is restricted to producing coherent states, in contrast to the single-qubit states required by the original protocol. A related procedure for this scenario we explored in Chapter 2 and the information was encoded in the polarization. This encoding the information remained essentially unitarily equivalent to the original single-qubit encoding. The question whether UBQC is possible when the client uses phase-encoded coherent states (where the unitary equivalence no longer holds) remains open. The approaches developed in this section have been applied to the task of amplifying coherent states *truly perfectly*, which can be achieved probabilistically when the number of possible phases is finite. This is the subject of the next section of this chapter.

10.2 Truly noiseless amplification of light

10.2.1 Introduction

There has recently been widespread theoretical and experimental interest in schemes for "noiseless" amplification of coherent states [123, 124, 141, 125, 126, 127, 142]. These schemes aim to implement the operation $|\alpha\rangle \rightarrow |g\alpha\rangle$, for $g > 1$ and any α. This is not possible to achieve perfectly with unit probability, but can be done probabilistically with arbitrarily high fidelity. Noiseless amplification could for example be used in quantum repeaters, or for entanglement purification through "breeding" larger Schrödinger cat states from "kittens", by probabilistically transforming $N_{\pm}(|\alpha\rangle \pm |-\alpha\rangle)$ into $N'_{\pm}(|g\alpha\rangle \pm |-g\alpha\rangle)$ with high fidelity.

Common to all existing schemes is that the amplification is not truly noiseless, or perfect, for non-zero success probability. That is, the fidelity approaches unity only in the limit of vanishing success probability. This must in fact hold for any phase-independent amplification scheme [126]. The suggested schemes achieve higher fidelity for smaller α or smaller gain, but it is only if either $|\alpha\rangle = |0\rangle$ or $g = 1$ that the fidelity can be 100% for non-zero success probability, in which case of course no amplification actually takes place. For experimental realisations, the overall success probability is usually not even quoted, and only the fidelity in case of successful operation is reported as a figure of merit. A complete and fair comparison of the different schemes is therefore difficult. The success probability, especially for schemes that involve single-photon states as resources, is nevertheless usually very low.

In this section we want to point out that in contrast to existing theoretical and experimental schemes, there *is* in fact a way to achieve truly noiseless amplification, that is, 100% fidelity, also for finite non-zero success probability and finite non-zero coherent state amplitudes. This is possible if one relaxes the demand that the amplification should work for any $|\alpha\rangle$, and instead selects any finite number of coherent states that one wants to amplify perfectly. The restriction to a finite set of states need not be serious, since many quantum information and communication protocols use a selected set of states, including quantum cryptography [86, 136, 137], blind quantum computing [1], and quantum digital signatures using coherent states, which was the main topic of PartII of this thesis. For example, the set of symmetric coherent states $|\alpha e^{im2\pi/N}\rangle$, where α is fixed and $m = 1, 2, \ldots N$, may be amplified truly perfectly with non-zero success

probability.

In fact, any set of linearly independent quantum states, coherent or other, may be amplified or cloned perfectly with a finite non-zero probability of success. This follows from the fact that linearly independent states may be unambiguously distinguished from each other with finite success probability [143]. Perfectly identifying a quantum state clearly allows us to fabricate an unlimited number of copies, or equivalently, to prepare a state with the same phase and arbitrarily high amplitude. Hence, it is not only possible to perfectly amplify any linearly independent set of states, but the average gain of truly noiseless probabilistic amplification can be *arbitrarily high*, since the success probability times the gain is unlimited. Moreover, unambiguous state discrimination of coherent states may be realized using only linear optics and non-photon-number-resolving photodetectors, without using auxiliary non-classical states [144, 145]. The same resources allow also realization of perfect amplification based on state discrimination.

If we do not require arbitrarily high gain, then the success probability can be higher than for schemes based on unambiguous state discrimination. For amplification of symmetric sets of coherent states, results on transforms between sets of symmetric states [8], presented in 10.1, are key to working out what processes are possible. Such transforms might be termed "umbrella transforms", if we visualize the symmetric states as the spines of an umbrella in a space of suitable dimensionality. A probabilistic transform that decreases pairwise overlaps, one example being noiseless amplification, may then be thought of as "opening the umbrella".

The remainder of this chapter is organized as follows. In Section 10.2.2, we briefly review unambiguous discrimination of coherent states using linear optics, and discuss how to use this for truly noiseless amplification. Definitions related to transformations between sets of quantum states are given in Section 10.2.3. In Section 10.2.4, we investigate truly noiseless amplification of coherent states, for finite gain, by viewing it as a transform between symmetric sets of states. As already mentioned, the success probability can then be higher than for procedures that use state discrimiation. It turns out that there are two regimes; small amplitude amplification, where the amplitudes of both initial and amplified states are below one, and a general regime where the amplitude of the final states, or of both initial and final states, are above one. As shown in Section 10.1, transforms between sets of states may be "leaky" or "leakless", depending on whether there is an extra "leak" state correlated with the desired output in the case of success. It turns out that in the small amplitude regime, the optimal "umbrella transforms" for noiseless amplification are leakless, whereas in the general regime, they may be leaky. We finish with a discussion.

10.2.2 Amplification of coherent states using linear optics

Ivanovic [128], Dieks [129] and Peres [130] realized that two non-orthogonal quantum states can be unambiguously distinguished from each other with a certain probability. That is, if the measurement succeeds, the result is always correct, but there is a chance that the measurement fails, giving an inconclusive result. The failure probability for the optimal procedure is equal to the overlap between the two quantum states. In the completely general case, optimal unambiguous measurements are not easy to find analytically [146, 147], but such a measurement

is at least possible as soon as at least one of the quantum states is linearly independent of the others [143].

For two coherent states $|\alpha\rangle$ and $|-\alpha\rangle$, the optimal measurement may be realized using only linear optics [144]. The state to be identified, $|\pm\alpha\rangle$, is directed onto a balanced beam splitter, with a fixed state $|\alpha\rangle$ incident on the other input port. If the phase relationships between output and input ports are arranged so that the beam splitter transforms $|\alpha\rangle_1 \otimes |\beta\rangle_2$ to $|(\alpha+\beta)/\sqrt{2}\rangle_1 \otimes |(\alpha-\beta)/\sqrt{2}\rangle_2$, we see that if the state to be identified was $|\alpha\rangle$, then output port 1 will contain $|\sqrt{2}\alpha\rangle$ and port 2 will be empty, and if it was $|-\alpha\rangle$, then output port 1 will be empty and output mode 2 contain $|-\sqrt{2}\alpha\rangle$. By detecting photons in the output ports, we can therefore unambiguously tell whether the state in input port 1 was $|\alpha\rangle$ or $|-\alpha\rangle$. Since any coherent state contains a vacuum component, we may not see any photons at all, which corresponds to the inconclusive outcome. The probability for this is $\langle 0|\sqrt{2}\alpha\rangle = \langle -\alpha|\alpha\rangle = \exp(-|\alpha|^2)$, which is the optimal (minimal) failure probability. Clearly, no photon counting is required, only being able to tell the difference between the vacuum and any nonzero number of photons.

For a balanced beam splitter with other phase relationships, we can adjust the phase of the fixed state in input port 2 so that the procedure still works. Also, if the two states to be distinguished are not $|\pm\alpha\rangle$ but $|\alpha\rangle$ and $|\beta\rangle$, then we can precede the described measurement with displacement of the unknown input mode, containing either the state $|\alpha\rangle$ or $|\beta\rangle$, by $-(\alpha+\beta)/2$ using a beam splitter, and then distinguish $|\pm(\alpha-\beta)/2\rangle$ using the technique above.

This unambiguous measurement may be used for perfect amplification as shown in Figure 10.5, where the first box shows a suggested way to prepare the states to be distinguished, and the second box the unambiguous measurement itself. The fact that we need to specify the phases of $|\pm\alpha\rangle$ implies that there exists a phase reference beam, which we without loss of generality assume to be $|\beta\rangle$, where α and β have the same phase, but different amplitude; a strong reference beam would have $|\beta| \gg |\alpha|$. The fixed state $|\alpha\rangle$ in input mode 2 is likely also split off this reference beam, as shown in Figure 10.5. Conditional on whether the state is identified as $|\alpha\rangle$ or $|-\alpha\rangle$, we implement the corresponding phase shift on the reference beam, giving the amplified state. The gain is then only limited by how strong the reference beam is. Alternatively, we could amplify relative to some other reference beam, not necessarily with the same phase as $|\alpha\rangle$ (but we still need the fixed state $|\alpha\rangle$ with the correct phase in input mode 2 for the unambiguous measurement).

A similar procedure is possible for distinguishing between more than two coherent states using linear optics, but will then not attain the optimal success probability [145]. In short, if there are N possible different states, then we can split the unknown state in N beams using a multiport, and test each component against one of the possible states (with amplitude suitably scaled down) using a beam splitter similar to as described above for two coherent states. If we manage to rule out all but one of the possible states, then we have unambiguously identified the input state as the remaining one. (Actually, we would only need to split the state to be identified in $N-1$ components, since if we manage to rule out all but one of the possible states, then the state must have been the remaining one.) The success probability of this procedure is non-zero, but not optimal. It can be somewhat improved by splitting the original state in M copies, with $M \to \infty$,

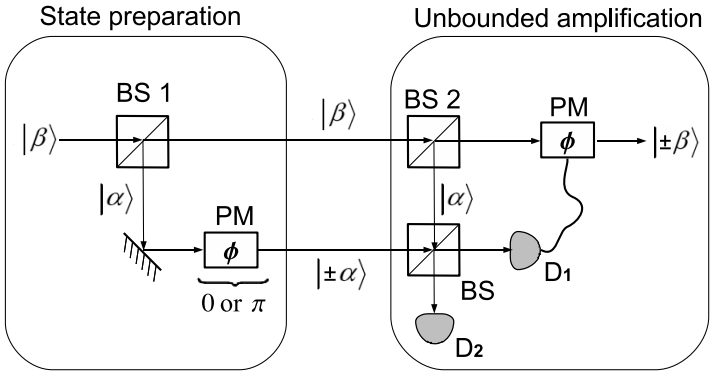

Figure 10.5. Truly perfect amplification of the states $|\pm\alpha\rangle$ *based on unambiguous discrimination, where we assume that* $|\beta| \gg |\alpha|$*. The beamsplitters denoted* BS 1 *and* BS 2 *split off a minor fraction of the strong beam* $|\beta\rangle$ *of amplitude of norm* $|\alpha|$*. The beamsplitter* BS *is balanced, and boxes labeled* PM *denote phase modulators. The amplification procedure fails only if both detectors* D_1 *and* D_2 *fail to detect a photon.*

still using only linear optics [145].

Any such procedure to unambiguously distinguish a finite number of coherent states may be used to noiselessly amplify them with a finite success probability and gain only limited by the strength of a reference beam, similar to amplification of two coherent states illustrated in Figure 10.5. If we manage to identify the state, we implement the corresponding phase shift on the reference beam. Although this requires only linear optics and detectors that do not resolve photon numbers, the disadvantage is that such a procedure cannot be used to amplify a superposition of the possible incident states. This obviously limits the usefulness when the superposition is important, such as when "breeding" larger cat states in order to enhance entanglement.

It is nevertheless in principle possible to realize truly noiseless amplification in such a way that superpositions are preserved. This is because one can in principle realize unambiguous state discrimination in two steps. First, one probabilistically transforms the selected set of non-orthogonal states into orthogonal ones without destroying possible superpositions of the states in the set, and this is followed by a measurement to distinguish the different orthogonal states. Truly noiseless amplification that preserves superpositions can then be achieved by omitting the final measurement, and only registering whether the first step succeeded or failed (how this works is also clarified by Eq. 10.2.3 in the following Section).

If the states to be amplified are symmetric to start with, then it follows from results of the last section on symmetric states transforms, that the success probability can be made independent of the initial state, and therefore the weights of the states in the superposition will be preserved. If the set of states is not symmetric, then the success probabilities for different states in the "source" set may not be equal. Amplification that preserves the superposition but re-weights the individual states is then still possible. Also, if we base the procedure on unambiguous state discrimination, then the amplified states in the superposition will be orthogonal, corresponding to infinite coherent state amplitude (the unambiguous measurement, if completed, would give us

perfect knowledge about which state was prepared, if only one of the initial states was prepared). Alternatively, if the amplitude of the amplified states is below one, then amplification that preserves superpositions is also possible, since then a leakless transform is possible, as we show in Sec. 10.2.4.1 (the concept of leak is introduced in the next section). We leave it open whether superposition-preserving truly noiseless amplification of coherent states could be realized using only linear optics.

Alternatively, we could remove the detectors and the strong reference beam $|\beta\rangle$ in the second box in Figure 10.5, and view the state exiting from the beam splitter labelled "BS", after combining the incident state $|\pm\alpha\rangle$ with the fixed state $|\alpha\rangle$, as an "amplification", with gain $\sqrt{2}$, of the incident state. This "amplification" occurs with unit probability, that is, it is deterministic, and transforms a superposition $N_{\pm}(|\alpha\rangle_1 \pm |-\alpha\rangle_1)$ into $N'_{\pm}(|\sqrt{2}\alpha\rangle_1 \otimes |0\rangle_2 \pm |0\rangle_1 \otimes |-\sqrt{2}\alpha\rangle_2)$. The deterministically "amplified" state will then be a superposition of different output modes. However, the overlap between the incident states $|\pm\alpha\rangle_1 \otimes |\alpha\rangle$ is necessarily the same as the overlap between the states $|\sqrt{2}\alpha\rangle_1 \otimes |0\rangle_2$ and $|0\rangle_1 \otimes |-\sqrt{2}\alpha\rangle_2$. Thus it is questionable if this process really could be called amplification, without subsequently combining the amplified states into the same spatial mode. That can only be done probabilistically, since otherwise we would be able to deterministically decrease the overlap of two quantum states, which is impossible.

Amplification with in principle unlimited gain will necessarily have the same optimal success probability as unambiguous state discrimination. We will now proceed to investigate the optimal success probabilities and other properties of the amplification transforms for specified finite levels of gain. For this, we first need to state some definitions related to transforms between sets of quantum states.

10.2.3 Transforms between sets of states

We consider two sets of N pure states, called the *source* and *target* sets, denoted (respectively)

$$A = \{|a_i\rangle\};\ B = \{|b_i\rangle\},$$

and a heralded probabilistic transform $\mathcal{T}$ which for input state $|a_i\rangle$ produces the state $|b_i\rangle$ with probability p_i, and a $|Fail\rangle$ state with probability $1-p_i$. By Theorem 3 in [133] such a transform exists iff there exists a unitary transform U performing

$$U|a_i\rangle = \sqrt{p_i}|b_i\rangle|\psi_i\rangle|0\rangle + \sqrt{1-p_i}|Fail\rangle|\phi_i\rangle|1\rangle,\ \forall i \tag{10.28}$$

for some sets of states $L = \{|\psi_i\rangle\}_N$ and $R = \{|\phi_i\rangle\}_N$. The states $|0\rangle$ and $|1\rangle$ are orthogonal. To complete the realization of $\mathcal{T}$, after the application of U the third register is measured in this basis, and, optionally, the second register may be traced out.

When the transform succeeds, the state $|b_i\rangle$ is generated along with a state $|\psi_i\rangle$, possibly correlated with the input state. This state leaks additional information about i, hence the set L is called

the *leak*. When the transform fails, the constant state $|Fail\rangle$ is produced along with the state $|\phi_i\rangle$ which may be correlated with the source state, and may be used to attempt a reconstruction of the target state $|b_i\rangle$. This set of states R we call the *redundancy*. The leak (redundancy) states are not correlated with the input state iff the states in the leak (redundancy) are identical for all source states, up to global phase. If the success probabilities do not depend on the input state, the transform is called *uniform*. For uniform transforms (of success probability p) the criterion (10.2.3) may be rewritten, in terms of the Gram matrices of the sets A,B, L and R respectively, as

$$G_A = p\, G_B \circ G_L + (1-p)\, G_R, \tag{10.28}$$

where $\circ$ denotes the Hadamard (point-wise) matrix product. The Gram matrix of a set of states $\{|c_i\rangle\}$ is defined as the square matrix with elements $\langle c_i|c_j\rangle$.

Finally, a finite set of states is *symmetric* if there exists a fixed unitary which, when applied on the i^{th} state, produces the $(i+1 \mod N)^{th}$ state. Symmetric states are interesting as they often appear in quantum protocols (e.g. many quantum key distribution schemes [86, 136, 137] and in blind quantum computing [1] and quantum digital signature schemes with coherent states, which was the topic of Part II of this thesis.

10.2.4 Amplification as state transforms

If the source set of coherent states we wish to amplify perfectly is known, and the required gain $g > 1$ is pre-set, then the amplification procedure becomes a particular type of state transform which we have presented in Section 10.1. Here, we will assume that the source set of coherent states is a symmetric set of N states. The source and target states are then

$$A = \left\{|a_i\rangle := |e^{i\theta_k}\alpha\rangle\right\}_{k=0}^{N-1}, B = \left\{|b_i\rangle := |e^{i\theta_k}\beta\rangle\right\}_{k=0}^{N-1}, \tag{10.29}$$

where $\theta_k = 2k\pi i/N$ and $\beta = g\alpha$. An amplification transform takes states from set A to corresponding (amplified) states in set B, and without the loss of generality we define the amplitudes α and β to be real positive numbers. The question is with what success probability such amplification is possible.

Since the set A is a set of linearly independent states, using state discrimination one can always perform a measure-and-prepare procedure, and in fact reach any desired, unlimited gain. Thus, the lower bound on the success of an amplification procedure is given by d_A, denoting the success probability of unambiguous state discrimination of the states in A. If we also take into the account the probability of unambiguous discrimination of states in B, an upper bound of the success probability of amplification can be derived. If d_A and d_B are the respective probabilities of optimal unambiguous discrimination of states in the sets A and B, then the corresponding amplification transform cannot succeed with a probability higher than

$$p_{up} = \frac{d_A}{d_B} \tag{10.30}$$

as a higher success probability would violate the optimality of d_A. Similar methods have been used to bound the success probability of decreasing the overlap of two quantum states, which includes state-dependent cloning or two states [132]. Similarly, one could derive other bounds by observing the optimal probabilities of minimum error measurements [113] on the sets A and B, or in fact any measurement optimizing any other figure of merit (e.g. maximal mutual information, maximum likelihood, etc.).

As we will show, the bound p_{up} can in fact be reached for source and target amplitudes below one, whereas for target state amplitudes above one it cannot always be saturated. The techniques we use have been developed in [134, 133, 8]. By the results given in [134, 133], an amplification transform succeeding with probability p exists if the equality given in equation (10.2.3) is satisfied for some Gram matrices of states [9] G_L and G_R, and where G_A and G_B are the Gram matrices of the source and amplified coherent states, respectively. Since A and B are symmetric sets of states, the matrices G_A and G_B are circulant [10], and hence diagonalize in the unitary discrete Fourier transform basis which is given by the columns or rows of the unitary discrete Fourier matrix of appropriate size N,

$$uDFT = 1/\sqrt{N}\left[\exp\ \frac{-2\,(p-1)\,(q-1)\,i\pi}{N}\right]_{p,q}. \tag{10.31}$$

Moreover, by Lemma 33, if there exists any amplification procedure for symmetric states succeeding with some success probability, then there exists an amplification procedure succeeding with the same success probability, where the leak and redundancy are symmetric sets of states.

Thus, in order to find the optimal success probability, we may assume that all the matrices appearing in the existence criterion (10.2.3) are circulant, and they all diagonalize in the unitary discrete Fourier transform basis. Criterion (10.2.3) may then be written in terms of vectors containing the eigenvalues of the Gram matrices as

$$\lambda_A = p\lambda_B * \lambda_L + (1-p)\lambda_R. \tag{10.32}$$

The vector λ_X above contains the diagonal elements of the matrix $uDFT^{\dagger}.G_X.uDFT$ which is diagonal when X is a symmetric set of states, and $*$ denotes the circular convolution of vectors, defined component-wise as

$$(\lambda_B * \lambda_L)_i = \frac{1}{N}\sum_{j=0}^{N-1}(\lambda_B)_j(\lambda_L)_{[(N-j+i)\bmod N]}. \tag{10.33}$$

For more details on the construction above see 10.1.

[9] A matrix is a Gram matrix of states if and only if it has unity across the diagonal and is positive semi-definite, see [133].

[10] A circulant matrix is a square matrix, whose i^{th} row is the right cyclic shift of the $(i-1)^{st}$ row.

All the results we will give rely on the properties of the spectrum of Gram matrices of coherent states which we give collectively in the Section 10.3 for convenience. As this spectrum has roughly two regimes of behaviour, depending on the amplitudes α and β being below or above one, we will separately address two distinct cases: small amplitude amplification (where $0 < \alpha \leq \beta \leq 1$), and general amplification (all other amplitude combinations). We begin by considering the scenario where both input and output amplitudes are small, i.e. less than one. From a practical standpoint, low amplitude amplification is of high importance since weak coherent states are often used in quantum information protocols. For sufficiently high amplitudes (also depending on N, that is, how many the states are), the symmetric sets of coherent states are effectively *classical*, that is, mutually almost orthogonal, and can be reliably distinguished. From a theoretical viewpoint, adhering to low amplitudes allows us to derive useful properties which do not hold for higher amplitudes.

10.2.4.1 Small amplitude amplification

If the amplitudes α and β of sets of symmetric coherent states A and B, respectively, satisfy $|\alpha| < |\beta| < 1$, the following two properties hold for the spectra of their corresponding Gram matrices G_A and G_B.

Property 1: the eigenvalues of G_A appear in strictly decreasing order, where the order is induced by the order of the diagonal elements of the diagonalized matrix obtained by the conjugation of G_A with the $uDFT$ matrix (cf. Lemma 39 below). This does not hold for higher amplitudes.

Property 2: the quotient of the last eigenvalues of G_A and G_B is smaller than the quotient of any other two corresponding eigenvalues (cf. Lemma 40 below and the derivation preceding it). Again, this holds only in the small amplitude regime.

For proofs, please see Section 10.3.

Property 2 above implies that the upper bound on the optimal success probability p_{up} in the low-amplitude regime, addressed in the beginning of this section, is reached in the leakless scenario, as we now show. First, we note the link between the optimal success probability d_S of uniformly and unambiguously discriminating a set of pure states S and the spectrum of the Gram matrix G_S of S: the optimal success probability d_S is equal to the smallest eigenvalue of G_S (this is easily derived from the results in [133, 3] as was done in the last section). Also, the sufficient criterion (10.32) for the existence of a probabilistic leakless transform taking the states from A to B where both sets of states are symmetric, succeeding with the probability p, can be written as

$$\lambda_A - p\lambda_B \geq 0, \tag{10.34}$$

where λ_A and λ_B are the vectors of eigenvalues of matrices G_A and G_B as discussed in the previous section. To see this, note that if the transform is leakless, then $\lambda_B * \lambda_L = \lambda_B$. The maximal possible p is then equal to $\min_j(\lambda_A^j/\lambda_B^j)$, where λ_A^j and λ_B^j are the j^{th} components

of the vectors λ_A and λ_B, respectively. Now, by the second property, this minimum is attained for the last eigenvalues (i.e. $j = N-1$), which is exactly the upper bound p_{up}. Thus, there exists a leakless transform saturating the upper bound of the success probability of amplification p_{up}.

Moreover, it can be shown by using **Property 1** that this bound is saturated *only* by a leakless transform in the small amplitude regime. From criterion (10.32), if there exists an amplification transform with a non-trivial leak, succeeding with some probability p, then the relation

$$\lambda_A - p\lambda_B * \lambda_L \geq 0 \tag{10.35}$$

holds, where λ_L is the vector of eigenvalues of the Gram matrix of the leak. Note that here we are assuming that the Gram matrix of the leak diagonalizes in the unitary discrete Fourier transform basis, which is justified without the loss of generality due to Lemma 33. If the leak is not trivial (not a fixed state) then λ_L is a vector of non-negative numbers adding up to N, at least two of which are not zero. Then note that the vector $\lambda_C = \lambda_B * \lambda_L$ contains the (normalized) weighted sums of the components of λ_B, the weights being the elements of λ_L (see the definition of the discrete convolution of vectors in expression (10.33)). Since the smallest component λ_B^{min} is the unique last component of λ_B (for $|\beta| < 1$ by **Property 1**), and at least two of the elements of λ_L are non-zero, the last component of λ_C is strictly greater than λ_B^{min}. But then it holds that

$$p \leq \frac{\lambda_A^{N-1}}{\lambda_C^{N-1}} < \frac{\lambda_A^{N-1}}{\lambda_B^{min}} = \frac{\lambda_A^{min}}{\lambda_B^{min}}. \tag{10.36}$$

Hence, the success probability of any leaky (non-leakless) amplification transform for low amplitudes is strictly less than optimal.

Thus we have shown that small amplitude amplification can be done optimally, i.e. saturating the obvious upper bound of the success probability p_{up}, and that this optimal transform is always leakless. The amplification procedure properties change significantly when one is interested in amplification involving states with amplitudes above unity, as we will see next.

10.2.4.2 General amplification

For 'any amplitude' amplification, i.e. when $\beta > 1$, we no longer have the convenient properties given in the previous subsection. In particular, optimal transforms can be leaky, in which case the upper bound p_{up} derived through the probabilities of unambiguous discrimination (see expression (10.30)) sometimes no longer can be reached. More formally, we have the following lemma:

Lemma 38. *Let λ_B^{min} be the smallest eigenvalue of the Gram matrix of the target, amplitude amplified, symmetric set of coherent states. Then if λ_B^{min} is a unique smallest eigenvalue then an optimal amplification transform with a non-trivial leak does not saturate the upper bound p_{up}.*

Proof:

Let c_j denote the j^{th} component of the vector $\lambda_C = \lambda_B * \lambda_L$, where λ_B and λ_L are vectors of eigenvalues of the Gram matrices of the target states and the leak states. Let $c^{min} = \min_j c_j$. Then, if λ_B^{min} is unique, and since $\lambda_B * \lambda_L$ contains the (normalized) weighted sums of the components of λ_B, the weights being the elements of λ_L, and at least two weights are not zero, it holds that $\lambda_B^{min} < c^{min}$.

Let p be the success probability of an optimal amplitude amplification transform with the leak characterised by λ_L. Then it holds that

$$\lambda_A - p\lambda_C \geq 0. \tag{10.37}$$

Also, due to optimality, for some component j it holds that

$$\lambda_A^j - pc_j = 0. \tag{10.38}$$

Assume first that $\lambda_A^j = \lambda_A^{min}$. Then

$$p = \frac{\lambda_A^{min}}{c_j}, \tag{10.39}$$

and because $\lambda_B^{min} < c_{min}$ it holds that

$$p = \frac{\lambda_A^{min}}{c_j} < \frac{\lambda_A^{min}}{\lambda_B^{min}} = p_{up}, \tag{10.40}$$

so the upper bound is not saturated.

Assume now that $\lambda_A^j \neq \lambda_{min}^A = \lambda_A^l$ for some position $l \neq j$.

Since

$$\lambda_A - p\lambda_C \geq 0 \tag{10.41}$$

it holds that

$$p \leq \min_i \frac{\lambda_A^i}{c_i}, \tag{10.42}$$

so since

$$p = \frac{\lambda_A^j}{c_j} \tag{10.43}$$

it holds that

$$p = \frac{\lambda_A^j}{c_j} \leq \frac{\lambda_A^l}{c_l} = \frac{\lambda_A^{min}}{c_l} \leq \frac{\lambda_A^{min}}{c_{min}} < \frac{\lambda_A^{min}}{\lambda_B^{min}} = p_{up}. \tag{10.44}$$

Therefore the upper bound is not attained, and the lemma holds. □

With Lemma 38 in place, we now show through an example that in the case of general amplification, the leakless case may not be optimal, and the upper bound p_{up} can sometimes no longer be obtained. Consider amplification of a symmetric set of 4 coherent states from amplitude $\alpha = 2$ to amplitude $\beta = 2.3$. The eigenvalues of the corresponding Gram matrices are then given by

$$\lambda_A = [0.976392, 0.971942, 1.02428, 1.02739]^T \tag{10.45}$$

$$\lambda_B = [1.00553, 0.991527, 0.99452, 1.00842]^T \tag{10.46}$$

and the upper bound is given with $p_{up} = 0.980248$. Note that the smallest eigenvalue of the Gram matrix of the target states is unique, so Lemma 38 can be applied, and the upper bound cannot be reached in the leaky setting.

What remains to be seen is what the success probability of a leakless transform is. The leak of a leakless transform are kets with only global phases possibly differing. Lemma 33 can still be applied in this case, hence we may assume that this leak is symmetric. This implies that the argument of the global phase of the k^{th} ket is "symmetric" as well and will be of the form $\theta_k = \pi kj/2$ for $j = 0, \ldots, 3$. By the properties of the discrete Fourier transform of powers of roots of unity, the vector of eigenvalues of such a leak will be a vector with all components zero, except at the position $((4 - j \mod 4) + 1$ where its entry is 4.

A convolution of a vector comprising zeroes, except at one position where the entry is one (or a constant c), with any other vector induces a circular permutation of the other vector (multiplied by the constant c). Hence, we can directly check the optimal leakless success probability of the leakless amplification procedure, by going through all the circular permutations of λ_B. We find that the optimal leakless transform succeeds with probability $p_{leakless} = 0.977298 < p_{up}$. So, the upper bound cannot be reached for the leakless scenario either, which means that, surprisingly, it cannot be reached at all. We note that although the values used in this analysis are numerical, the discrepancies the conclusion relies on (i.e. the uniqueness of the smallest eigenvalue and comparison of magnitude of the quotients) are well within numerical precision, hence the conclusion is unlikely to be a numerical artifact.

Now, is there a leaky transform that does not saturate the bound, but does better than the best leakless transform? Using the optimization technique developed in 10.1 we find that the success probability of an optimal transform for this example is $p_{opt} = 0.978604$ which is slightly larger than the optimal leakless transform $p_{leakless} = 0.977298$, and, necessarily, strictly below the upper bound $p = 0.980248$.

To summarize, we have proven the following:

- The success probability of amplifying a symmetric set A of N coherent states of amplitude α to the states in a symmetric set B of coherent states of a larger amplitude β, for small amplitudes $|\alpha| < |\beta| < 1$, can reach the upper bound imposed by the ratio of success probabilities of optimal unambiguous discrimination of sets A and B, respectively.
- For small amplitudes $|\alpha| < |\beta| < 1$ the optimal transform is always leakless.

- The optimal success probability of amplification of small amplitudes is explicitly given by

$$p_{opt} = \frac{\sum_{r=0}^{\infty} \frac{\alpha^{2(N(r+1)-1)}}{(N(r+1)-1)!}}{\sum_{r=0}^{\infty} \frac{\beta^{2(N(r+1)-1)}}{(N(r+1)-1)!}}. \tag{10.47}$$

  (please see the equation (10.68) in Section 10.3, and the subsequent paragraph).

- If $|\beta| > 1$, the numerical testing we have performed indicates that the upper bound imposed by the ratio of the success probabilities for unambiguous discrimination of the states in sets A and B cannot always be reached, and optimal transforms may be leaky.

10.2.5 Conclusions

In this work, we have shown that truly noiseless amplification of coherent states is possible if one only requires the amplification to work perfectly for a finite number of of states. Similarly, perfect cloning of any other linearly independent states is also possible, and amplification is clearly closely related to cloning. Depending on whether the amplitude of the amplified "target" states are below or above one, the optimal success probability may be simply obtained, or require optimization techniques like the ones we have presented in 10.1. The average gain is in principle unlimited, since it is possible to base the amplification on unambiguous state discrimination. In case of success, this allows us to prepare an amplified state with arbitrary high amplitude. If we require a finite level of gain, the optimal success probability is higher than for unlimited gain. We have also explained how to implement truly noiseless amplification based on unambiguous state discrimination using only linear optics.

If we visualize the N coherent states to be amplified as the spines in an umbrella in an N-dimensional space, then noiseless amplification of these states, which decreases their pairwise overlaps, may be thought of as "opening the umbrella". Sometimes the optimal amplification procedures may result in extra "leak" and "redundancy" states, apart from the desired amplified states. The leak and the redundancy may be correlated with and therefore carry information about the input state. Since the optimal "umbrella transform" for truly noiseless amplification is always leakless when the amplitude of the amplified (target) states is below one, as we have shown, this regime may be convenient if cryptographic aspects come into consideration. For example, in a two-party protocol, where Alice sends some quantum states to Bob who is supposed to further transform them, Alice can monitor the success probability declared by Bob. If it is optimal, she knows that there can be no additional leak (assuming that Alice uses some other way of checking that when Bob does declare that the process has succeeded, he has indeed obtained the quantum state he is supposed to). A related situation arises in blind quantum computing, where Alice wants to run a quantum computation on Bob's quantum computer without Bob learning about her data or her algorithm [1]. In the original scheme, Alice is required to prepare single-qubit states. If Alice only can prepare, say, weak coherent states, then one possibility may be for Alice to require Bob to turn these into single-qubit states in such a way that Alice can monitor any additional information Bob may gain. Such transforms from symmetric coherent states to

symmetric qubit states were considered in Section 10.1. If the amplitude of the target states is above one, then the optimal "umbrella" amplification transform may be leaky.

A few years ago, quantum cloning attracted widespread attention, see e.g. [148, 149, 150]. Amplification and cloning are closely connected, especially for coherent states, since for example the state $|\alpha\rangle \otimes |\alpha\rangle$ may be transformed into $|\sqrt{2}\alpha\rangle$ using a beam splitter, and vice versa. More generally, if $g = \sqrt{N}$, then the state $|g\alpha\rangle$ is equivalent to N copies of $|\alpha\rangle$, in the sense that $|\sqrt{N}\alpha\rangle$ can be transformed into N copies of $|\alpha\rangle$ (and vice versa) by a linear optical network (a balanced multiport). It is well known that perfect universal quantum cloning, i.e. of arbitrary states, is impossible [148, 150]. On the other hand, probabilistic perfect cloning of linearly independent states is possible [122]. This mirrors the fact that probabilistic perfect amplification of linearly independent states is possible with finite success probability.

To elaborate on the connection to cloning, the existing schemes for "noiseless" probabilistic amplification of coherent states are (almost) perfect cloners for coherent states, but do not clone superpositions of coherent states as well. For example, choosing $g = \sqrt{2}$, if $|\alpha\rangle \to |\sqrt{2}\alpha\rangle$ for any α, then the "cat" state $N_{\pm}(|\alpha\rangle \pm |-\alpha\rangle)$ would change into $N'_{\pm}(|\sqrt{2}\alpha\rangle \pm |-\sqrt{2}\alpha\rangle)$, which may be transformed into $N'_{\pm}(|\alpha\rangle \otimes |\alpha\rangle \pm |-\alpha\rangle \otimes |-\alpha\rangle)$ using a balanced beam splitter. This state is not equal to $N_{\pm}(|\alpha\rangle \pm |-\alpha\rangle) \otimes N_{\pm}(|\alpha\rangle \pm |-\alpha\rangle)$, that is, to two copies of the original cat state. This is similar to the simple proof that universal cloning is impossible [148].

Nevertheless, this feature is not a disadvantage when perfect amplification schemes are used e.g. to enhance entanglement. That the operation $|\alpha\rangle \to |g\alpha\rangle$ only has unit fidelity in the limit of vanishing success probability, on the other hand, is a disadvantage. If we select a finite linearly independent set of states $|\alpha_i\rangle$ which the amplification should work perfectly for, then the fidelity of the probabilistic process $N \sum_i c_i|\alpha_i\rangle \to N' \sum_i c_i|g\alpha_i\rangle$ can be truly perfect, as we have pointed out. The price we have to pay is that the scheme *must* be dependent on phase and amplitude. Nevertheless, since such schemes can be realized with only linear optics, as discussed in Sec. 10.2.2, we expect them to be of great interest for quantum information applications.

10.3 Technical results

Here we give the proofs of the Lemmas and other technical results which were states and used throughout this chapted. For the reasons of brevity, occasionally we will skip through technical details, and rather present the main ideas. We begin by proving the uniformization (Lemma 31) and symmetrization (Lemma 33) lemmas.

10.3.1 Proof of Lemmas 31 and 33

Lemma 31 *(Uniformization) If there exists a probabilistic transformation $\mathcal{T}$ taking the states in A to states in B which succeeds with the probabilities $\{p_i\}_{i=1}^N$, where A and B are symmetric sets of states, then there exists a uniform probabilistic transform $\mathcal{T}'$ taking the states in A to*

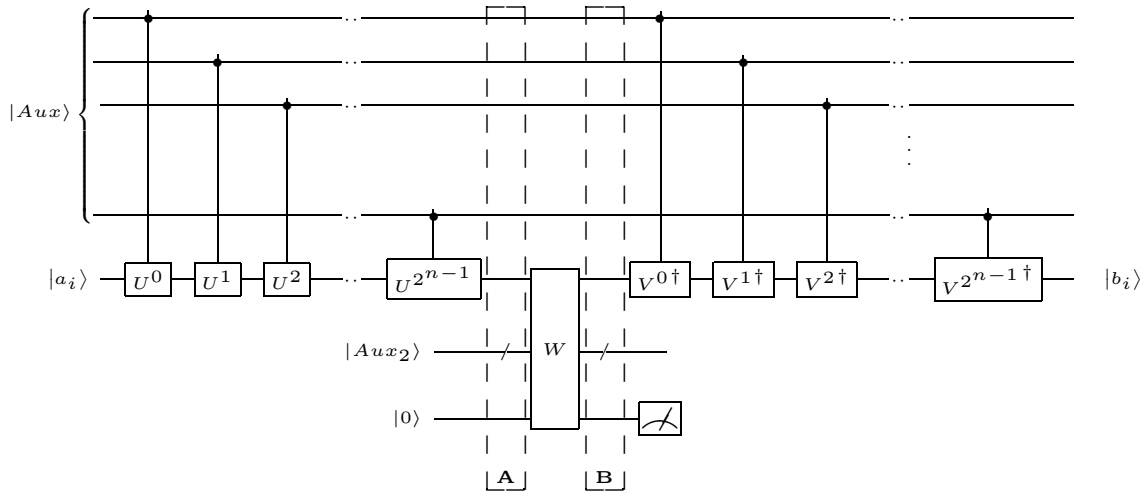

*Figure 10.6. The quantum circuit by which a probabilistic transform, realized by the action of a unitary W acting on an augmented Hilbert space followed by the measurement of an indicator register, can be 'uniformized'. The same circuit also serves to symmetrize the leak and the redundancy of a uniform probabilistic transform. In the proofs of lemmas we will address the states of the system above at cuts **A** and **B** denoted in this figure.*

states in B which succeeds with probability

$$p = \frac{1}{N}\sum_{i=1}^{N} p_i.$$

Proof:

As noted, each probabilistic transform may be realized as a unitary transform acting on an augmented Hilbert space, followed by a measurement of an indicator register. In the circuit of Figure 10.6, the transform $\mathcal{T}$ is represented by this extended unitary W. As both the input and output sets of states are symmetric, there exist unitaries which sequentially shift though the states of the set, obeying the intrinsic order. We denote these unitaries by U and V, corresponding to the sets A and B respectively, and the controlled powers of these unitaries appear in the circuit. The state $|Aux\rangle$ is pre-set to be the uniform superposition

$$|Aux\rangle = 1/\sqrt{N}\sum_{k=0}^{N-1}|k\rangle, \tag{10.47}$$

where $|k\rangle$ is the l qubit state of the computational basis $|b_{l-1}\rangle \otimes \cdots \otimes |b_0\rangle$, $b_j \in \{0,1\}$ for all j such that $(b_{l-1}\ldots b_0)_2 = (k)_{10}$, where the subscripts designate the base of the number representations.

First let us show that the circuit shown performs the desired transform. The state of the system

at cut **A** in the circuit is

$$1/\sqrt{N}\sum_{k=0}^{N-1}|k\rangle|a_{i+k\ mod\ N}\rangle|Aux_2\rangle|0\rangle, \tag{10.48}$$

where $|Aux_2\rangle$ is some fixed auxiliary state in a sufficiently dimensional state space. The notation we shall use corresponds to the notation used in formula 10.1. Following this, the transform $\mathcal{T}$ is applied to the register which contained the input state. The transform is explicitly realized as a unitary W acting on a bigger space. The state at cut **B** in the circuit is

$$\begin{aligned}1/\sqrt{N}\left(\sum_{k=0}^{N-1}\sqrt{p_{i+k\mod N}}|k\rangle|b_{i+k\mod N}\rangle|\psi_{i+k\mod N}\rangle\right)|0\rangle \\ +1/\sqrt{N}\left(\sum_{k=0}^{N-1}\sqrt{1-p_{i+k\mod N}}|k\rangle|Fail\rangle|\phi_{i+k\ mod\ N}\rangle\right)|1\rangle.\end{aligned} \tag{10.48}$$

If the measurement outcome of the indicator (the last) register corresponds to the state $|0\rangle$, then the transform has succeeded (c.f. expression 10.1). From the expression above, it can be seen that this happens with probability $p = \frac{1}{N}\sum_{i=1}^{N} p_i$. Assume that the indicator measurement yielded the desired output. The section of the circuit after cut B undoes the controlled rotations, and the state at the end of the entire circuit is

$$\mathcal{N}\sum_{k=0}^{N-1}\sqrt{p_{i+k\mod N}}|k\rangle|b_i\rangle|\psi_{i+k\mod N}\rangle. \tag{10.49}$$

The middle register contains the desired output state, and the rest of the system contains a new leak. This overall procedure constitutes the new, uniformized probabilistic transform $\mathcal{T}'$ from the statement of the Lemma, which succeeds with the averaged probability p. This proves Lemma 31. □

One can verify that the new leak, generated by the 'uniformized' transform described above, comprises a symmetric set of states. Using an analogous analysis, one can show that the redundancy (state generated in case of the measurement outcome corresponding to the $|1\rangle$ state in the indicator register) is a symmetric set as well. Now, if the extended unitary W corresponds to a uniform probabilistic transform, with leak and redundancy which are not symmetric, then the extended transform of Figure 10.6 will have the same success probability as W itself, and the leak and redundancy will be symmetrized. Thus, the analysis above proves Lemma 33 as well. □

Lemma 32 *A Gram matrix of kets is a circulant matrix if and only if the corresponding set of kets is symmetric.*

Proof:
Let $A = \{|a_k\rangle\}_{k=0}^{N-1}$ be a set of kets. We first show the necessity. If the set of kets is symmetric, then its Gram matrix is circulant. Let U be the unitary which sequentially shifts through the set

of kets, obeying the intrinsic order. Then the Gram matrix may be written as

$$G_A = [\langle a_p|a_q\rangle]_{p=0,q=0}^{N-1,N-1} = \left[\langle a_0|U^{\dagger p}U^q|a_0\rangle\right]_{p=0,q=0}^{N-1,N-1} = \left[\langle a_0|U^{q-p}|a_0\rangle\right]_{p=0,q=0}^{N-1,N-1} = \left[\langle a_0|U^{q-p \mod N}|a_0\rangle\right]_{p=0,q=0}^{N-1,N-1}. \quad (10.49)$$

It is easy to verify that the last matrix in the sequence of equalities above is circulant.

Next we show the sufficiency. If A is a set of states such that its Gram matrix G_A is circulant, then it is symmetric. Since G_A is a matrix of states, we have that G_A allows the Cholesky decomposition, that its spectrum $\{\lambda_k\}_{k=0}^{N-1}$ is real, non-negative and sums up to N (as the trace is preserved under basis change), and as it is circulant, we have that it diagonalizes in the $uDFT$ basis. Using these properties and a bit of matrix algebra, one can show that if a set of kets $\{|\psi_k\rangle\}_{k=0}^{N-1}$ has G_A as a Gram matrix, then its elements can be written as

$$|\psi_k\rangle = \frac{1}{\sqrt{N}}\sum_{j=0}^{N-1}\frac{1}{\sqrt{\lambda_j}}e^{\frac{2kj\pi i}{N}}|b_j\rangle, \quad (10.50)$$

where the kets $\{|b_k\rangle\}_{k=0}^{N-1}$ comprise an orthonormal basis, and we define the coefficient $\frac{1}{\sqrt{\lambda_j}}$ to be zero if $\lambda_j = 0$. Consider the unitary U, acting on the $\{|b_j\rangle\}_j$ basis as follows:

$$U|b_j\rangle = e^{\frac{2j\pi i}{N}}|b_j\rangle. \quad (10.51)$$

By applying U on the ket $|\psi_k\rangle$ we have:

$$U|\psi_k\rangle = \frac{1}{\sqrt{N}}\sum_{j=0}^{N-1}\frac{1}{\sqrt{\lambda_j}}e^{\frac{2kj\pi i}{N}}U|b_j\rangle = \frac{1}{\sqrt{N}}\sum_{j=0}^{N-1}\frac{1}{\sqrt{\lambda_j}}e^{\frac{2(k+1)j\pi i}{N}}|b_j\rangle = |\psi_{k+1 \mod N}\rangle. \quad (10.51)$$

Hence, the set of kets A which they represent is symmetric and this proves the lemma. □

The following lemmas were given in Section 10.1.6. In their proofs we shall adhere to the notation of that section.

Lemma 37 *Let the amplitude α of the states in the set A, defined in equation (10.29), satisfy $0 < \alpha \leq 1$. Then there exists a uniform multiprobabilistic transform with the success probability vector $(p_0, \ldots, p_{N-1})$, which takes the states from the set A to the collection of target states $\{B^j\}_{j=1}^{N-1}$* and *is redundancy-free and leakless. The failure probability p_0 of this transform is equal to* $\exp(-2\alpha^2)$.

Proof:
As noted, the desired transform exists if and only if

$$G_A = p_0\Pi^f + p_1 G_{B^1}\circ\Pi^1 + \cdots + p_{k-1}G_{B^{N-1}}\circ\Pi^{N-1} \quad (10.52)$$

holds for a vector of probabilities $(p_0, \ldots, p_{N-1})$ and for a a set of Gram matrices of states $\{\Pi^f, \Pi^1, \ldots, \Pi^{N-1}\}$. Acknowledging the requirement that this transform is leakless and redundancy-free the criterion becomes

$$G_A = p_0 \mathbf{1} + p_1 G_{B^1} + \cdots + p_{N-1} G_{B^{N-1}}, \tag{10.53}$$

where $\mathbf{1}$ is a matrix with all entries being the unity.

The matrix G_{B^j} can be written as

$$G_{B^j} = \underbrace{G_B \circ \cdots \circ G_B}_{j \text{ times}} := G_B^{\circ j}, \tag{10.54}$$

and since G_A, G_B are circulant, and the Hadamard product of circulant matrices is circulant, G_{B^j} is circulant for all j. Hence all the matrices in expression (10.53) simultaneously diagonalize in the unitary discrete Fourier transform basis so we can write this criterion in terms of vectors of eigenvalues of the corresponding matrices:

$$\lambda_{G_A} = p_0 \lambda_{\mathbf{1}} + p_1 \lambda_{G_{B^1}} + \cdots + p_{N-1} \lambda_{G_{B^{N-1}}}. \tag{10.55}$$

The vector $\lambda_{\mathbf{1}}$ is the first vector of the canonical basis, that is vector with one as the first entry and zeroes elsewhere, multiplied by N.

It can be shown that, for any N, the vector of eigenvalues of G_B has only the first two eigenvalues non-zero, and their value is $N/2$. From this, using the properties given in expressions (10.7) and (10.8), we can see that, for $k \leq N-1$, the vector of eigenvalues of $G_B^{\circ k}$ is given by

$$\lambda_{B^k} = \frac{N}{2^k} \left[\binom{k}{0}, \binom{k}{1}, \cdots, \binom{k}{k}, 0, \cdots, 0 \right]^T. \tag{10.56}$$

Let M be the column matrix defined by

$$M = \left[N e_1 | \lambda_B | \lambda_{B^2} | \cdots | \lambda_{B^{N-1}} \right]. \tag{10.57}$$

Then we can rewrite the condition (10.55) as a system of equations,

$$\lambda_{G_A} = M \overrightarrow{p} \tag{10.58}$$

where $\overrightarrow{p} = [p_0, \ldots, p_{N-1}]^T$. Since M is upper-triangular, with non-zero element across the diagonal, it is invertible. Hence, there exists a unique vector $\overrightarrow{p}$ satisfying the system above. The sum of the elements of a column of the matrix M is N, so we can see (by multiplying the system (10.58) with the row vector $\frac{1}{N}[1, \ldots, 1]$ from the left) that $\sum_{i=0}^{N-1} p_i = 1$, as the sum of the eigenvalues of G_A is N.

To prove the stated Lemma, we need to show that all the values p_i are non-negative (for $0 < \alpha \leq 1$), and that we need to show that $p_0 = \exp(-2\alpha^2)$. We begin by showing the positivity of

values p_i, as stated, and we finish of the proof by showing that $p_0 = \exp(-2\alpha^2)$.

As noted above, the system (10.58) has a unique solution (and $M^{(-1)}$ exists), and we need to show that the solution vector comprises positive elements, i.e.

$$M^{(-1)}\lambda_{G_A} \tag{10.59}$$

is a vector of non-negative real numbers. Note that the matrix M can be written as $M = M'.D$, where M' collects all the binomial coefficients and D is a diagonal matrix which appropriately assigns the weights to the columns of M. The k^{th} column of matrix M' is then given by

$$\left[\binom{k}{0}, \binom{k}{1}, \cdots, \binom{k}{k}, 0, \cdots, 0\right]^T.$$

The inverse of M is then

$$M^{(-1)} = D^{-1}.M'^{(-1)}. \tag{10.59}$$

As the matrix D^{-1} comprises only positive elements (moreover it is also diagonal), in order to show that the expression (10.59) is a non-negative vector, it will suffice to show that

$$M'^{(-1)}\lambda_{G_A} \tag{10.60}$$

is a non-negative vector. Let S be a diagonal matrix of size N of alternating signs, the first sign being positive. Using known properties of sums of binomial coefficients, one can show that $S.M'.S$ is the inverse of the matrix M'. We omit the proof of this claim as the proof is technical, and the details are of no further consequence.

Now we proceed to show that each entry of the vector

$$M'^{(-1)}\lambda_{G_A} = S.M'.S\lambda_{G_A} \tag{10.61}$$

is non-negative, if the amplitude α is a positive and less or equal to unity. Let λ_i be the i^{th} eigenvalue of the matrix G_A, *i.e.* the i^{th} component of λ_{G_A}. Note that the enumeration starts at zero. Then the k^{th} entry of the vector $S.M'.S\lambda_{G_A}$ is given by

$$(e_k)^T S.M'.S\lambda_{G_A} = \sum_{j=k}^{N-1}(-1)^{j+k}\binom{j}{k}\lambda_j \tag{10.62}$$

The last entry of the vector $S.M'.S\lambda_{G_A}$ is the last eigenvalue of G_A, hence positive, so for the expression (10.62) to be positive, it suffices to show that

$$\binom{j}{k}\lambda_j - \binom{j+1}{k}\lambda_{j+1} \geq 0 \tag{10.63}$$

for all $0 \geq k < N-1$ and $k \leq j < N-1$. This expression simplifies to

$$\binom{j}{k}\lambda_j - \binom{j+1}{k}\lambda_{j+1} = \binom{j}{k}\left(\lambda_j - \frac{j+1}{j-k+1}\lambda_{j+1}\right). \tag{10.64}$$

Since $\binom{j}{k}$ is positive, we only need to show that the following holds:

$$\lambda_j - \frac{j+1}{j-k+1}\lambda_{j+1} \geq 0. \tag{10.65}$$

In order to show this, we need to analyse the structure of the eigenvalues appearing as components of λ_{G_A}. Recall, λ_{G_A} was defined as the discrete Fourier transform of the first row of G_A. Using the expansion of coherent states in the Fock basis the j^{th} eigenvalue can be given as

$$\lambda_j = \sum_{l=0}^{N-1} \exp\left(-2jl\pi i/N\right) \sum_{r=0}^{\infty} e^{-\alpha^2} \frac{\alpha^{2r}}{r!} \exp(2lr\pi i/n). \tag{10.66}$$

This can further be rearranged as follows:

$$\begin{aligned} \lambda_j &= e^{-\alpha^2} \sum_{l=0}^{N-1} \sum_{r=0}^{\infty} \exp\left(-2jl\pi i/N\right) \frac{\alpha^{2r}}{r!} \exp(2lr\pi i/n) && (10.67)\\ &= e^{-\alpha^2} \sum_{r=0}^{\infty} \frac{\alpha^{2r}}{r!} \sum_{l=0}^{N-1} \exp(2l(r-j)\pi i/n), && (10.68) \end{aligned}$$

where in order to get to expression (10.66), we used the fact that the infinite sum above is absolutely convergent, hence allows the commuting of sums.

By the properties of sums of roots of unity, the expression $\sum_{l=0}^{N-1} \exp(2l(r-j)\pi i/n)$ is equal to n if $r-j$ is divisible by N and zero otherwise. Hence we get

$$\lambda_j = e^{-\alpha^2} N \sum_{r=0}^{\infty} \frac{\alpha^{2(Nr+j)}}{(Nr+j)!}. \tag{10.69}$$

The elements in the sum above appear as the summands in the Taylor expansion of $e^{2\alpha}$; for $j = 0$, this sum collects every N^{th} summand from the Taylor series expansion, starting from the zeroth summand. For any other j it collects every N^{th} summand from the Taylor series expansion, starting from the j-th summand. We note that the eigenvalues above, for a fixed N can be expressed in a closed form in terms of Generalized hypergeometric functions.

We set out to show that inequality (10.65) holds. By inserting the explicit expressions for the eigenvalues we have derived, we obtain the expression

$$\lambda_j - \frac{j+1}{j-k+1}\lambda_{j+1} =$$
$$e^{-\alpha^2} N \left(\sum_{r=0}^{\infty} \frac{\alpha^{2(Nr+j)}}{(Nr+j)!} - \frac{j+1}{j-k+1} \sum_{r=0}^{\infty} \frac{\alpha^{2(Nr+j+1)}}{(Nr+j+1)!} \right), \tag{10.69}$$

and again by absolute convergence of the sums above we may reshuffle them and obtain

$$e^{-\alpha^2} N \sum_{r=0}^{\infty} \frac{\alpha^{2(Nr+j)}}{(Nr+j)!} \left(1 - \frac{j+1}{(j-k+1)} \frac{1}{(Nr+j+1)} \alpha^2\right). \tag{10.70}$$

The expression above is positive if the expression in the last parenthesis is positive. Now we inspect the coefficient with the term α^2 in the parenthesis,

$$\frac{j+1}{(j-k+1)} \frac{1}{(Nr+j+1)}.$$

This expression is always positive, and note that the denominator $(j-k+1)$ is greater or equal to unity, and the denominator $(Nr+j+1)$ is larger or equal to $j+1$, so the entire expression is less or equal to unity. But then for $\alpha \leq 1$ the expression (10.70) is non-negative.

To finish the proof we need to show that $p_0 = \exp(-2\alpha^2)$. Note that $p_0 = e_0^T M^{(-1)} \lambda_{G_A}$. Recall, $\lambda_{G_A} = DFT.(e_0^T.G_A)^T$ (i.e. the DFT of the first row of the Gram matrix of the set A is the vector of eigenvalues of G_A). The exact form of M^{-1} was given in expression (10.59), and we can see that

$$e_0^T M^{(-1)} = \frac{1}{N} [1, -1, 1, \ldots, 1, -1].$$

Thus, it holds that

$$p_0 = e_0^T M^{(-1)} DFT.(e_0^T.G_A)^T = \frac{1}{N} [1, -1, 1, \ldots, 1, -1] .DFT.(e_0^T.G_A).$$

We can see that that

$$[1, -1, 1, \ldots, 1, -1] .DFT = N e_{N/2},$$

as this is equivalent to adding a π phase to each of the rows of the DFT matrix and then summing up the rows. Without the phase shift, the sum of the rows is a vector with a non-zero entry only at the first position. The phase shift corresponds to a cyclic permutation of columns by $N/2-1$ positions, so the sum of the rows of the permuted DFT matrix has the only non-zero entry at the $(N/2+1)^{st}$, and this entry is N. Hence we have

$$p_0 = e_{N/2}.(e_0^T.G_A)^T = \exp(-2\alpha^2),$$

and we have proven our Lemma. □

10.3.2 Properties of the spectrum of the Gram matrix of symmetric sets of coherent states

The vector of eigenvalues of the Gram matrix of a symmetric set of coherent states λ_{G_A} can be obtained by the discrete Fourier transform of the first row of G_A (for details, see Section 10.1).

Hence, the j^{th} eigenvalue can be given as

$$\lambda_j = \sum_{l=0}^{N-1} \exp\left(-2jl\pi i/N\right) \langle \alpha | \alpha \exp\left(2l\pi i/N\right) \rangle. \tag{10.66}$$

Using the expansion of the coherent states in the Fock number basis the expression above can be written as

$$\lambda_j = \sum_{l=0}^{N-1} \exp\left(-\frac{2jl\pi i}{N}\right) \sum_{r=0}^{\infty} e^{-\alpha^2} \frac{\alpha^{2r}}{r!} \exp\left(\frac{2lr\pi i}{N}\right) \tag{10.67}$$

This can further be rearranged as follows:

$$\begin{aligned} \lambda_j &= e^{-\alpha^2} \sum_{l=0}^{N-1} \sum_{r=0}^{\infty} \exp\left(-2jl\pi i/N\right) \frac{\alpha^{2r}}{r!} \exp\left(2lr\pi i/N\right) \\ &= e^{-\alpha^2} \sum_{r=0}^{\infty} \frac{\alpha^{2r}}{r!} \sum_{l=0}^{N-1} \exp\left(-2jl\pi i/N\right) \exp(2lr\pi i/n) \\ &= e^{-\alpha^2} \sum_{r=0}^{\infty} \frac{\alpha^{2r}}{r!} \sum_{l=0}^{N-1} \exp(2l(r-j)\pi i/N), \end{aligned} \tag{10.66}$$

where to get to the the step (10.66) we used the fact that the infinite sum is absolutely convergent, thereby allowing the commuting of sums.

By the properties of sums of roots of unity, the expression $\sum_{l=0}^{N-1} \exp(2l(r-j)\pi i/n)$ is equal to n if $r-j$ is divisible by N and zero otherwise. Hence we obtain

$$\lambda_j = e^{-\alpha^2} N \sum_{r=0}^{\infty} \frac{\alpha^{2(Nr+j)}}{(Nr+j)!}. \tag{10.67}$$

The elements in the sum above appear as the summands in the Taylor expansion of $e^{2\alpha}$. For any j this sum collects every N^{th} summand from the Taylor series expansion starting from the j^{th} summand. We note that the eigenvalues above can be expressed in a closed form in terms of generalized hypergeometric functions. Using the presented form of the eigenvalues λ_j we can show that for amplitudes below unity, the order of eigenvalues is monotonously decreasing:

Lemma 39. *Let A be the symmetric set of N coherent states as defined in expression (10.29). Let λ_A be the vector of eigenvalues of the Gram matrix G_A generated by taking the discrete Fourier transform of the first row of G_A. If λ_j is the j^{th} component of λ_A, then for the real amplitude $\alpha \leq 1$ the eigenvalues in Λ are decreasingly ordered:*

$$\lambda_j \geq \lambda_{j+1}. \tag{10.68}$$

Proof:

We will show that $\lambda_j - \lambda_{j+1} \geq 0$. By using expression (10.67) derived above, we obtain

$$\begin{aligned} & \lambda_j - \lambda_{j+1} \\ = \; & e^{-\alpha^2} N \sum_{r=0}^{\infty} \frac{\alpha^{2(Nr+j)}}{(Nr+j)!} - e^{-\alpha^2} N \sum_{r=0}^{\infty} \frac{\alpha^{2(Nr+j+1)}}{(Nr+j+1)!} \\ = \; & e^{-\alpha^2} N \alpha^{2j} \sum_{r=0}^{\infty} \frac{\alpha^{2Nr}}{(Nr+j)!} \left(1 - \alpha^2 \frac{1}{Nr+j+1}\right), \end{aligned} \tag{10.69}$$

where the last step is possible due to absolute convergence of the sums above. Note that the expression above is positive if $\left(1 - \alpha^2 \frac{1}{Nr+j+1}\right)$ is positive. It holds that $Nr + j + 1 \geq 1$, so for $\alpha \leq 1$ the expression above is positive and we have our claim. Note also that in the case where α is strictly less than unity and positive, λ_j is strictly greater than λ_{j+i}. So, for amplitudes below 1, the probability of success of unambiguous discrimination of symmetric sets of coherent states is given by the last eigenvalue in the vector λ_A. This eigenvalue is given by

$$\lambda_{min} = e^{-\alpha^2} N \sum_{r=0}^{\infty} \frac{\alpha^{2(N(r+1)-1)}}{(N(r+1)-1)!}. \tag{10.68}$$

In the case **Property 2** holds, from the equation above we can give the explicit optimal success probability of amplification of a set of symmetric coherent states. This is simply the quotient of the respective values of λ_{min} for the two amplitudes, in the low amplitude regime.

In the remainder of this section we prove **Property 2** from the main body of text. Let $\lambda_j(\alpha)$ be the j^{th} eigenvalue of the Gram matrix of the symmetric set of N coherent states of (real) amplitude α. **Property 2** states that

$$\frac{\lambda_j(\alpha)}{\lambda_j(\beta)} \geq \frac{\lambda_{N-1}(\alpha)}{\lambda_{N-1}(\beta)} \tag{10.69}$$

for all $j = 0, \ldots, N-1$, and $0 < \alpha < \beta < 1$. Since all the eigenvalues are positive and non-zero, the inequality above can be rewritten as

$$\frac{\lambda_j(\alpha)}{\lambda_{N-1}(\alpha)} \geq \frac{\lambda_j(\beta)}{\lambda_{N-1}(\beta)} \tag{10.70}$$

which holds iff $\lambda_j(x)/\lambda_{N-1}(x)$ is a decreasing function on $(0, 1)$. Note that the functions $\lambda_j(x)$ are non-negative for all j on the interval of interest. If it is the case that $\lambda_j(x)/\lambda_{j+1}(x)$ is a decreasing function on the interval $(0, 1)$ for all $j = 0, \ldots, N-2$, then the function $\lambda_j(x)/\lambda_{N-1}(x)$ is decreasing as well, which would imply **Property 2**. To see this, note that the equality

$$\frac{\lambda_j(x)}{\lambda_{j+1}(x)} \frac{\lambda_{j+1}(x)}{\lambda_{j+2}(x)} \cdots \frac{\lambda_{N-2}(x)}{\lambda_{N-1}(x)} = \frac{\lambda_j(x)}{\lambda_{N-1}(x)} \tag{10.71}$$

holds for every j, and since the left-hand side of the expression above is a product of positive decreasing functions, the right-hand side must also be a decreasing function. Hence, it will suffice

to show that $\lambda_j(x)/\lambda_{j+1}(x)$ is a decreasing function on the interval of interest, which we state as the following Lemma.

Lemma 40. *The quotient of eigenvalues*

$$\frac{\lambda_j(x)}{\lambda_{j+1}(x)} \tag{10.72}$$

is a decreasing function on $(0,1)$ *for all* $j = 0, \ldots, N-2$.

Proof:

By recalling the analytic expression for the eigenvalues, given in (10.67), we have

$$\begin{aligned} \frac{\lambda_j(x)}{\lambda_{j+1}(x)} &= \frac{e^{-x^2} N \sum_{r=0}^{\infty} \frac{x^{2(Nr+j)}}{(Nr+j)!}}{e^{-x^2} N \sum_{r=0}^{\infty} \frac{x^{2(Nr+j+1)}}{(Nr+j+1)!}} \\ &= \frac{\sum_{r=0}^{\infty} \frac{x^{2(Nr+j)}}{(Nr+j)!}}{\sum_{r=0}^{\infty} \frac{x^{2(Nr+j+1)}}{(Nr+j+1)!}}. \end{aligned} \tag{10.72}$$

Let us introduce the notation

$$l_j(x) = \sum_{r=0}^{\infty} \frac{x^{2(Nr+j)}}{(Nr+j)!}. \tag{10.73}$$

To prove Lemma 40 we then need to show that $l_j(x)/l_{j+1}(x)$ is a decreasing function on $(0,1)$ for all $j = 0, \ldots, N-2$. Note that the functions $l_j(x)$ are positive, strictly increasing and infinitely differentiable functions. Also, using the same technique we applied to prove the analogous property for the eigenvalues themselves, it holds that $l_j(x) \geq l_{j+1}(x)$ for all $j = 0, \ldots N-2$, and for $x \in (0,1)$. Then, the quotient $l_j(x)/l_{j+1}(x)$ is decreasing in x if and only if the derivative of the quotient over x is non-positive on the interval of interest:

$$\frac{l_j'(x) l_{j+1}(x) - l_j(x) l_{j+1}'(x)}{(l_{j+1}(x))^2} \leq 0 \tag{10.74}$$

Since the denominator of the fraction above is always positive, this inequality holds if and only if the inequality

$$l_j'(x) l_{j+1}(x) - l_j(x) l_{j+1}'(x) \leq 0 \tag{10.75}$$

holds.

It is easy to verify the following property of the derivatives of the functions $l_j(x)$:

$$l_j'(x) = \frac{d}{dx} l_j(x) = 2x l_{j-1 \bmod N}(x). \tag{10.76}$$

Hence we have

$$\begin{aligned} &l_j'(x)l_{j+1}(x) - l_j(x)l_{j+1}'(x) \\ &= 2x\left(l_{j-1\,mod\,N}(x)l_{j+1}(x) - l_j(x)l_j(x)\right) \end{aligned} \tag{10.76}$$

which is non-positive on the interval of interest if and only if

$$l_{j-1\,mod\,N}(x)l_{j+1}(x) - l_j(x)l_j(x) \leq 0. \tag{10.77}$$

Note that if $j = 0$ the expression above resolves to

$$l_{N-1}(x)l_1(x) - l_0(x)l_0(x) \leq 0. \tag{10.78}$$

Since $l_{N-1}(x) \leq l_0(x)$ and $l_1(x) \leq l_0(x)$ on the interval $(0, 1)$, and since all the values these functions attain are positive, we have that for $j = 0$ the condition given in expression (10.77) holds. By using the definitions of the functions $l_j(x)$, for $j = 1, \ldots N-2$, we obtain

$$\begin{aligned} l_{j-1}(x)l_{j+1}(x) - l_j(x)l_j(x) = &\sum_{r=0}^{\infty} \frac{x^{2(Nr+j-1)}}{(Nr+j-1)!} \sum_{r=0}^{\infty} \frac{x^{2(Nr+j+1)}}{(Nr+j+1)!} - \\ &\sum_{r=0}^{\infty} \frac{x^{2(Nr+j)}}{(Nr+j)!} \sum_{r=0}^{\infty} \frac{x^{2(Nr+j)}}{(Nr+j)!} = \\ x^{4j} &\sum_{r=0}^{\infty} (x^{2N})^r \frac{1}{(Nr+j-1)!} \sum_{r=0}^{\infty} (x^{2N})^r \frac{1}{(Nr+j+1)!} - \\ &x^{4j} \sum_{r=0}^{\infty} (x^{2N})^r \frac{1}{(Nr+j)!} \sum_{r=0}^{\infty} (x^{2N})^r \frac{1}{(Nr+j)!}. \end{aligned} \tag{10.78}$$

The sign of the expression above is then equal to the sign of the expression

$$\begin{aligned} &\sum_{r=0}^{\infty} (x^{2N})^r \frac{1}{(Nr+j-1)!} \sum_{r=0}^{\infty} (x^{2N})^r \frac{1}{(Nr+j+1)!} - \\ &\sum_{r=0}^{\infty} (x^{2N})^r \frac{1}{(Nr+j)!} \sum_{r=0}^{\infty} (x^{2N})^r \frac{1}{(Nr+j)!}. \end{aligned} \tag{10.78}$$

Note that to prove that $l_{j-1}(x)l_{j+1}(x) - l_j(x)l_j(x) \leq 0$ for $j > 0$ (and consequently **Property 2**), it will suffice that the expression (10.78) is negative for all $x \in (0, 1)$, and for $j = 1, \ldots N-2$. Also, since any positive power is a bijection on the interval $x \in (0, 1)$, and we require negativity on the entire interval, the expression(10.78) is negative if and only if the expression

$$\begin{aligned} &\sum_{r=0}^{\infty} (x^{N})^r \frac{1}{(Nr+j-1)!} \sum_{r=0}^{\infty} (x^{N})^r \frac{1}{(Nr+j+1)!} - \\ &\sum_{r=0}^{\infty} (x^{N})^r \frac{1}{(Nr+j)!} \sum_{r=0}^{\infty} (x^{N})^r \frac{1}{(Nr+j)!} \end{aligned} \tag{10.78}$$

is negative on the same interval.

Consider now the family of functions

$$f_j(x) = \sum_{r=0}^{\infty} \frac{x^{(Nr+j)}}{(Nr+j)!}.$$

Using the same construction as for the functions $l_j(x)$, it is easy to see that $f_j(x)/f_{j+1}(x)$ is a decreasing function on $(0,1)$ for $j = 1, \ldots N-2$ if and only if the expression (10.78) is negative on the same interval. For these functions f_j it is also easy to see that they are positive, strictly increasing, infinitely differentiable, and $f_j(x) \geq f_{j+1}(x)$ holds on the interval of interest for $j = 0, \ldots N-2$. It also holds that

$$\frac{d}{dx} f_j(x) = f_{j-1 \, mod \, N}(x). \tag{10.78}$$

Recall the property of log-concavity: a function is log-concave (on an interval) if the logarithm of that function is concave on the same interval. For functions which are twice differentiable, log-concavity holds if and only if the quotient of the derivative of the function and the function itself is decreasing (on the same interval). Hence, the requirement that $f_j(x)/f_{j+1}(x)$ is decreasing on the interval of interest is equivalent to the requirement that $f_{j+1}(x)$ is a log-concave function.

Here we invoke the following result given in Lemma 1 of the manuscript [151], also a consequence of the Lemma 3 in the Appendix of [152] (a published version of the aforementioned manuscript):

Lemma 41. *Let $g(x)$ be a strictly monotonic, twice differentiable function on the interval (a,b). Let also $g(a) = 0$ or $g(b) = 0$. Then if the derivative $g'(x)$ is log-concave on the same interval, $g(x)$ is log-concave on the interval.*

Since for all $j > 0$ the function $f_j(0)$ is zero, and all the functions f_j are strictly increasing, it holds that $f_{j+1}(x)$ is log-concave if $f_j(x)$ is log-concave. Inductively, if $f_1(x)$ is log-concave, so is $f_j(x)$ for all $j = 2, \ldots, N-1$. To finish the proof of Lemma 40 and thus of **Property 2**, we finally need to show that $f_1(x)$ is log-concave on $(0,1)$. Recall that $f_1(x)$ is log-concave on the interval of interest if the quotient $f_0(x)/f_1(x)$ is decreasing on the interval. This holds if the inequality

$$\begin{aligned} &f_0'(x)f_1(x) - f_1'(x)f_0(x) \\ &= f_{N-1}(x)f_1(x) - f_0(x)f_0(x) \leq 0 \end{aligned}$$

holds. But since we have that $f_j(x) \geq f_{j+1}(x)$ holds on the interval of interest for $j = 0, \ldots N-2$ and since all the functions above attain positive values, this inequality is satisfied. Hence Lemma 40 and **Property 2** are proven.

We note that the functions $l_j(x)$ and $f_j(x)$ are sub-series of the Taylor expansion of the functions e^{x^2} and e^x about the point $x = 0$, respectively, and as such are absolutely convergent, which

allows the unrestricted reshuffling of sums.

Definitions of the relevant classical and quantum complexity classes

Throughout this thesis we have been referring to various computational complexity classes, which we briefly define here. In general, a *computational complexity class* is a set of *problems* characterised by the resources which are required to solve them. More often than not, the basic blueprint of a definition of a complexity class is of the following form.

A computational complexity class C*, is the set of problems of instance size* n *which can be solved on an abstract machine* M *using* $O(f(n))$ *of a particular resource* $\mathtt{R}$*.*

For instance, the abstract machines can be Turing machines, boolean circuits, non-deterministic Turing machines, quantum circuits, etc. The only resources we will consider are *time* (number of computational steps), and *space*, both of which need to be defined for each abstract machine. Finally, the function order $O(f(n))$ is most often either "polynomial" or "exponential" the meaning of which is straightforward. The most common type of problems one considers are *decision problems*, which are the types of problems which have a binary "YES" or "NO" solution. A typical example of a decision problem is the general boolean satisfiability problem: given a boolean formula (of size n), determine whether there exists an assignment of the boolean variables which renders the entire formula true. For instance, the problem of factoring an integer is technically not a decision problem (as the output of this problem has to be a non-trivial factor of the input integer), but rather a *search* or an *optimization* problem. Nonetheless, many such decision problems admit a reformulation which is a decision problem – such a "reformulation" is called a *reduction* [11]. In this sense, while we will technically talk about complexity classes of decision problems we will be actually having in mind both decision problems and an other types of problems which are reducible to decision problems, such as factoring, the travelling salesman, linear programming, etc.

We begin with the smallest class mentioned in this thesis.

P is the set of decision problems in which both "YES" and "NO" answers can be given using a *deterministic Turing machine*, in a *polynomial* number of computational steps in the problem instance size. Instead of using a deterministic Turing machine in the definition, we could have used *uniform families of boolean circuits* comprising at most polynomially many gates in the input instance size. Thus, this class of problems can in practice be solved on standard computers we use every day.

[11] There are many ways to reformulate the factoring problem as a decision problem. For instance if one can solve the problem: "is the k^{th} binary digit of the smallest non-trivial factor of the input integer zero?", then by solving this problem a logarithmic number of times in the input size, one finds the factor.

NP is the set of decision problems in which "YES" instances can be decided on a *non-deterministic Turing machine*, in a *polynomial* number of computational steps in the problem instance size. This is a generalization of the P class, and has the following operational characterisation: it is the class of decision problems for which the "YES" instances have proofs of the fact that the answer is indeed "YES", which can be *verified* within the class P. For instance, both integer factoring and boolean satisfiability are in NP, and this is easily seen by using the second definition. In the case of integer factoring, if one is given a non-trivial factor of the input, this solution is easily verified in a polynomial number of steps on a deterministic Turing machine, by performing integer division of the input with the factor. For the boolean satisfiability problem, if one is given the satisfying sequence of variable values, one can easily check whether the sequence satisfies the formula (renders it true), as boolean formula evaluation can be done in polynomial time on a deterministic Turing machine. In the two examples we have given, the factor and the satisfying assignment are the "proofs" we have referred to in the characterisation of NP.

BPP is the set of decision problems in which both "YES" and "NO" answers can be given using a *probabilistic Turing machine*, in a *polynomial* number of computational steps in the problem instance size, with error probability of $1/3$.

Intuitively, a probabilistic Turing machine can be thought of as an algorithm which is allowed to toss a fair coin to make decisions throughout its run-time. If the input problem instance is a "YES" instance, the algorithm will output "YES" with probability above $2/3$. This property is often called *completeness*. If the correct answer is "NO" it will output "YES" with probability no higher than $1/3$. This property is called *soundness*.

The bounded error property ensures that by running the algorithm a number of times (say, at most a polynomial number of times) the probability of making the correct decision by taking the majority vote of the algorithm outputs approaches unity exponentially in the number of runs. In this sense the acceptable error bound need not be $2/3$ exactly, but in fact even the value $1/2 + 1/poly(n)$ will do, for any polynomial $poly(n)$ in the input instance size n. This is the class of problems that are considered tractable, or efficiently solvable in practice on classical computers.

ZPP is closely related to BPP. It is the set of decision problems in which both "YES" and "NO" answers can be given using a *probabilistic Turing machine*, in an *expected polynomial* number of computational steps in the problem instance size, with *zero error probability*. There are two differences between BPP and ZPP. First of all, the machine in ZPP class is not allowed to err. Secondly, the running time is required to be polynomial on average, however it is allowed to occasionally be much longer. Although the second difference may make us believe ZPP could be more powerful than BPP (while the first one may make us believe the opposite), it is known that $\mathsf{ZPP} \subseteq \mathsf{BPP}$.

BQP is a quantized version of BPP. This is the set of decision problems in which both "YES" and "NO" answers can be given using a *quantum Turing machine*, in a *polynomial* number of computational steps in the problem instance size, with error probability of $1/3$. Similarly to the case of classical classes, instead of talking about quantum Turing machines, one can talk about quantum circuits (of at most polynomial size) or measurement-based quantum computations (where the resource state, or total number of measurements is at most polynomial in the input instance size). This is the class of problems efficiently solvable on a quantum computer. It holds that $\mathsf{BPP} \subseteq \mathsf{BQP}$.

PH stands for the class of decision problems solvable within the polynomial hierarchy. The polynomial hierarchy is a hierarchy of decision problems generalizing the class P, NP and the complement of NP called $co-\mathsf{NP}$. The formal definition of this class in terms of the classes P and NP is technical and would require us to first introduce the concepts of *oracle machines*. For this reason we refrain from giving a formal definition, and direct the interested reader to [13]. For our purposes we will only state that the class PH can be defined as a countable union of "levels" of the polynomial hierarchy. The PH is said to collapse at the k^{th} level if all the higher levels of the hierarchy are contained in the k^{th}. For instance, if $\mathsf{P} = \mathsf{NP}$ then PH collapses to the zeroth level, and if $\mathsf{NP} = co-\mathsf{NP}$ then it collapses to the first level. The PH is believed not to collapse, and it contains all the classes mentioned so far, and is contained in the class PSPACE which we define next.

PSPACE is the set of decision problems in which both "YES" and "NO" answers can be given using a deterministic Turing machine, in a polynomial amount of *space* (on the Turing machine tape) in the problem instance size. In this class the resource of interest is *memory* rather than time. This class is believed to strictly contain all the classes we have mentioned so far.

EXPTIME is the set of decision problems in which both "YES" and "NO" answers can be given using a *deterministic Turing machine*, in an *exponential time* (number of computational steps) in the problem instance size.

NEXPTIME is the set of decision problems in which both "YES" and "NO" answers can be given using a *non-deterministic Turing machine*, in an *exponential time* (number of computational steps) in the problem instance size.

These last two classes are the exponential analogues of P and NP, respectively, and are real "heavyweights" of complexity theory. Both contain all the classes mentioned so far.

IP is a class of decision problems solvable by an interactive proof system. In interactive proof systems we consider a two-party setting, comprising a *verifier* and a *prover*. The prover is computationally unlimited (can solve all decision problems) and the verifier is a limited machine, capable of deciding problems in BPP only. The pair, the verifier and prover are presented with

a decision problem, and the verifier is allowed to communicate with the prover in an attempt to solve the problem. The number of rounds of communication is at most polynomial in the input instance size. At the end of the interaction, the verifier decides whether the solution to the problem is "YES" or "NO". Similarly to the other bounded-error classes a decision problem is solvable in IP if completeness and soundness are ensured: if the input problem instance is a "YES" instance, the verifier will output "YES" with probability above $2/3$, and if the correct answer is "NO" the verifier will output "YES" with probability no higher than $1/3$. Note that the prover can always decide any (decidable) problem, however it cannot be trusted. For instance, any BPP problem is in IP since the verifier can decide those problems on its own. Also, the IP class can be seen as a generalization of the NP class. As noted, the problems in NP have a "short proof" of the validity of the solution which can be verified by a P machine. Hence, to solve an NP problem, the verifier simply asks the prover for this short proof, and checks the declared solution. Thus, NP problems can be resolved with interactive proof systems with a single round of communication. If one allows for a polynomial number of communication rounds, even PSPACE problems can be decided by the verifier. Formally, we have $\mathsf{IP} = \mathsf{PSPACE}$.

QIP is the quantum analogue to the class IP, in which the verifier is "upgraded" from a BPP machine to a BQP machine. The messages sent comprise qubits instead of classical bits. While it has been shown that $\mathsf{QIP} = \mathsf{PSPACE}$ it has also been shown that $\mathsf{QIP}(3) = \mathsf{QIP}$, *i.e.* a quantum interactive proof system with a constant (3) rounds of communications has the same decision power as a quantum interactive proof system with polynomially many rounds. This stands in contrast to the classical case, where it is not known that $\mathsf{IP} = \mathsf{IP}(const)$ for any constant number of communication rounds.

MIP is a class of problems solvable by a multi-prover interactive proof system. Here, the verifier has access to two (or more) independent, computationally unbounded provers, and again has to decide a problem maintaining soundness and completeness like in the class IP. Crucially, the provers are not allowed to communicate in run-time. Intuitively, the advantage the verifier gets from access to two non-communicating provers resembles the advantages of cross-interrogation of criminals, using the information gained from one culprit against the other one. This seemingly drastically raises the decision powers of the verifier, and it has been proven that $\mathsf{MIP} = \mathsf{NEXPTIME}$.

QMIP is the quantized version of MIP where the verifier is a BQP machine, and the messages exchanged are allowed to be quantum messages. The servers are not allowed to communicate classically, but versions where the provers share no, a limited amount, or unlimited amount of entanglement, may alter the overall decision power of the system. If no entanglement is shared between the provers, then it is known that $\mathsf{QMIP} = \mathsf{MIP}$.

MIP^*. As we mentioned, in the QMIP setting, the deciding power of the quantum multi-prover interactive proof system may depend on the amount of entanglement the provers share. For this

reason, a setting where the verifier is a BPP machine, the provers cannot communicate, but are allowed to share entanglement. Recently in [79] it has been shown that QMIP $=$ MIP* using universal blind quantum computation, which we have briefly discussed in Chapter 4.

Completeness of a problem for a class Informally speaking, complete problems for a problem class C are the hardest problems in that class, where "hardness" of a problem is measured in terms of what one could compute, if she could compute that particular problem. More formally, a problem c in a decision problem class C is hard for the class C if there exists a many-to-one reduction (another decision problem) for any problem in C to c. This means that there exists an algorithm which "translates" any problem $c' \in C$ to an instance of the problem c, and the solution to this instance is the solution to the original problem c'. The "translation" – the many-to-one reduction – is usually required to be weaker than the class C itself, otherwise the entirety of the class C becomes hard for C. Any problem which is hard for a class C and is in C is said to be *complete* for C. Typical examples of complete problems are the complete problems for the NP class, such as the general boolean satisfiability problem, the travelling salesman, the subgraph isomophism problem, all of which are in NP, but also any other problem in NP is reducible to these problems under *polynomial reductions* (*i.e.* the reduction itself is in P). Not all classes have complete problems under interesting reductions, however, the class of problems solvable on a quantum computer BQP does. One example is the famous approximation of the Jones polynomial (the discovery which resulted from the study of anyonic topological quantum computation), but others exist such as the the Quadratically Signed Weight Enumerator problem, the Local Hamiltonian Eigenvalue Sampling problem, and the $Q - CIRCUIT$ problem [153, 19], all of which are complete for the class BQP under polynomial reductions. Often, the types of reductions we are interested may vary. For example, *Turing* reductions would allow an algorithm (solving $c' \in C$) to call a subroutine (solving the problem $c \in C$) many times (and it is clearly allowed to post-process the output of the subroutines), in contrast to the "strict reductions" we have assumed above. It is not known whether the classes of complete problems expand under these more lenient reductions. Also, the notion of completeness can also be defined for functional, optimization and sampling problems as well, as long as we maintain the spirit of the idea – that the ability to solve one (complete) problem, in some sense *easily* reduces to the ability to solve all problems in a given class.

References

[1] Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. Universal blind quantum computation. *Foundations of Computer Science, Annual IEEE Symposium on*, 0:517–526, 2009.

[2] Daniel Gottesman and Isaac Chuang. Quantum digital signatures. *arXiv:quant-ph/0105032v2*, 2001.

[3] Erika Andersson, Marcos Curty, and Igor Jex. Experimentally realizable quantum comparison of coherent states and its applications. *Phys. Rev. A*, 74(2):022304, Aug 2006.

[4] Vedran Dunjko, Elham Kashefi, and Anthony Leverrier. Blind quantum computing with weak coherent pulses. *Phys. Rev. Lett.*, 108:200502, May 2012.

[5] Elham Kashefi Tomoyuki Morimae, Vedran Dunjko. Ground state blind quantum computation on the aklt state. 2012. submitted for publication.

[6] P. J. Clarke, R. Collins, V. Dunjko, E. Andersson, J. Jeffers, and G. S. Buller. Experimental demonstration of quantum digital signatures. submitted for publication, 2012.

[7] Vedran Dunjko and Elham Kashefi. Algebraic characterisation of one-way patterns. In *DCM*, pages 85–100, 2010. Proceedings Sixth Workshop on Developments in Computational Models: Causality, Computation, and Physics.

[8] Vedran Dunjko and Erika Andersson. Transformations between symmetric sets of quantum states. *Journal of Physics A: Mathematical and Theoretical*, 45(36):365304, 2012.

[9] Vedran Dunjko and Erika Andersson. Truly noiseless probabilistic amplification. submitted for publication, 2012.

[10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information (Cambridge Series on Information and the Natural Sciences)*. Cambridge University Press, 1 edition, January 2004.

[11] Ulf Leonhardt. *Measuring the Quantum State of Light*. Cambridge University Press, 1997.

[12] Stephen M. Barnett. *Quantum Information*. Oxford University Press, 2009.

[13] Michael Sipser. *Introduction to the Theory of Computation*. Course Technology, 1 edition, December 1996.

[14] David Deutsch. Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 400(1818):97–117, 1985.

[15] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, June 1982.

[16] Andrew M. Childs. Secure assisted quantum computation, Jul 2005.

[17] Howard Barnum, Claude Crépeau, Daniel Gottesman, Adam Smith, and Alain Tapp. Authentication of quantum messages. In *Proceedings of the 43rd Symposium on Foundations of Computer Science*, FOCS '02, pages 449–458, Washington, DC, USA, 2002. IEEE Computer Society.

[18] P. Arrighi and L. Salvail. Blind quantum computation. *International Journal of Quantum Information*, 4:883–898, 2006.

[19] Dorit Aharonov, Michael Ben-Or, and Elad Eban. Interactive proofs for quantum computations. *Proceeding of ICS2010*, 2010.

[20] D. Aharonov and M. Ben-Or. Fault-tolerant quantum computation with constant error. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, STOC '97, pages 176–188, New York, NY, USA, 1997. ACM.

[21] R. Rivest, L. Adleman, and M. Dertouzos. On data banks and privacy homomorphisms. pages 169–177. Academic Press, 1978.

[22] C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st annual ACM symposium on Theory of Computing*, pages 169–178. ACM, 2009.

[23] M. Abadi, J. Feigenbaum, and J. Kilian. On hiding information from an oracle. *Journal of Computer and System Sciences*, 39:21–50, 1989.

[24] R. Raussendorf and H. J. Briegel. A one-way quantum computer. *Physical Review Letters*, 86:5188 – 5191, 2001.

[25] R. Raussendorf, D. E. Browne, and H. J. Briegel. Measurement-based quantum computation with cluster states. *Physical Review A*, 68:022312 [32 pages], 2003.

[26] V. Danos, E. Kashefi, and P. Panangaden. The measurement calculus. *Journal of ACM*, 54:8, 2007.

[27] R. Jozsa. An introduction to measurement based quantum computation. 2005. Available on arXiv:quant-ph/0508124v2.

[28] Charles H. Bennett and Stephen J. Wiesner. Communication via one- and two-particle operators on einstein-podolsky-rosen states. *Phys. Rev. Lett.*, 69:2881–2884, Nov 1992.

[29] Artur K. Ekert. Quantum cryptography based on bell's theorem. *Phys. Rev. Lett.*, 67:661–663, Aug 1991.

[30] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters. Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Phys. Rev. Lett.*, 70:1895–1899, Mar 1993.

[31] D. Gottesman and I. L. Chuang. Quantum teleportation is a universal computational primitive. *Nature*, 402, 1999.

[32] M. Hein, W. Dür, J. Eisert, R. Raussendorf, M. Van den Nest, and H. J. Briegel. Entanglement in Graph States and its Applications. February 2006.

[33] M. Hein, J. Eisert, and H. J. Briegel. Multiparty entanglement in graph states. *Phys. Rev. A*, 69:062311, Jun 2004.

[34] Vincent Danos, Elham Kashefi, and Prakash Panangaden. Parsimonious and robust realizations of unitary maps in the one-way model. *Phys. Rev. A*, 72:064301, Dec 2005.

[35] V. Danos and E. Kashefi. Determinism in the one-way model. *Physical Review A*, 74:052310 [6 pages], 2006.

[36] Daniel Gottesman. Theory of fault-tolerant quantum computation. *Phys. Rev. A*, 57:127–137, Jan 1998.

[37] Damian Markham and Barry C. Sanders. Graph states for quantum secret sharing. *Phys. Rev. A*, 78:042309, Oct 2008.

[38] Claude Crépeau, Daniel Gottesman, and Adam Smith. Secure multi-party quantum computation. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, STOC '02, pages 643–652, New York, NY, USA, 2002. ACM.

[39] D. Browne, E. Kashefi, M. Mhalla, and S. Perdrix. Generalized flow and determinism in measurement-based quantum computation. *New Journal of Physics*, 9, 2007.

[40] Einar Pius. Msc thesis: Automatic parallelisation of quantum circuits using the measurement based quantum computing model. http://www.epcc.ed.ac.uk/wp-content/uploads/2010/12/Einar%20Pius.pdf.

[41] Anne Broadbent and Elham Kashefi. Parallelizing quantum circuits. *Theor. Comput. Sci.*, 410(26):2489–2510, 2009.

[42] Elham Kashefi Joseph Fitzsimons. Unconditionally verifiable blind computation. arXiv:1203.5217v1, 2012.

[43] Andris Ambainis, Michele Mosca, Alain Tapp, and Ronald De Wolf. Private quantum channels. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 547–553. IEEE Computer Society, 2000.

[44] Stefanie Barz, Elham Kashefi, Anne Broadbent, Joseph F. Fitzsimons, Anton Zeilinger, and Philip Walther. Demonstration of Blind Quantum Computing. *Science*, 335(6066):303–308, January 2012.

[45] Valerio Scarani, Helle Bechmann-Pasquinucci, Nicolas J. Cerf, Miloslav Dušek, Norbert Lütkenhaus, and Momtchil Peev. The security of practical quantum key distribution. *Rev. Mod. Phys.*, 81:1301–1350, Sep 2009.

[46] N. Lütkenhaus and A. J. Shields. Focus on quantum cryptography: Theory and practice. *New Journal of Physics*, 11(4):045005, 2009.

[47] R Raussendorf, J Harrington, and K Goyal. Topological fault-tolerance in cluster state quantum computation. *New Journal of Physics*, 9(6):199, 2007.

[48] D. Gross and J. Eisert. Novel schemes for measurement-based quantum computation. *Phys. Rev. Lett.*, 98:220503, May 2007.

[49] Gavin K. Brennen and Akimasa Miyake. Measurement-based quantum computer in the gapped ground state of a two-body hamiltonian. *Phys. Rev. Lett.*, 101:010502, Jul 2008.

[50] Jianming Cai, Akimasa Miyake, Wolfgang Dür, and Hans J. Briegel. Universal quantum computer from a quantum magnet. *Phys. Rev. A*, 82:052309, Nov 2010.

[51] Akimasa Miyake. Quantum computational capability of a 2d valence bond solid phase. *Annals of Physics*, 326(7):1656 – 1671, 2011.

[52] Andrew S Darmawan, Gavin K Brennen, and Stephen D Bartlett. Measurement-based quantum computation in a two-dimensional phase of matter. *New Journal of Physics*, 14(1):013023, 2012.

[53] M. Van den Nest, K. Luttmer, W. Dür, and H. J. Briegel. Graph states as ground states of many-body spin-12 hamiltonians. *Phys. Rev. A*, 77:012301, Jan 2008.

[54] Tomoyuki Morimae and Keisuke Fujii. Blind topological measurement-based quantum computation. 2011.

[55] C. E. Shannon. A mathematical theory of communication. *Bell system technical journal*, 27, 1948.

[56] R. Renner. Security of quantum key distribution. PhD Thesis, 2005. http://arxiv.org/abs/quant-ph/0512258.

[57] Michael Ben-Or, MichałHorodecki, Debbie W. Leung, Dominic Mayers, and Jonathan Oppenheim. The universal composable security of quantum key distribution. In *Proceedings of the Second international conference on Theory of Cryptography*, TCC'05, pages 386–406, Berlin, Heidelberg, 2005. Springer-Verlag.

[58] Universally composable privacy amplification against quantum adversaries. In *TCC*, pages 407–425, 2005.

[59] Robert König, Renato Renner, Andor Bariska, and Ueli Maurer. Small accessible quantum information does not imply security. *Phys. Rev. Lett.*, 98(14):140502, Apr 2007.

[60] Christine Guerlin, Julien Bernu, Samuel Deleglise, Clement Sayrin, Sebastien Gleyzes, Stefan Kuhr, Michel Brune, Jean-Michel Raimond, and Serge Haroche. Progressive field-state collapse and quantum non-demolition photon counting. *Nature*, 448(7156):889–893, August 2007.

[61] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

[62] H.K. Lo, X. Ma, and K. Chen. Decoy state quantum key distribution. *Phys. Rev. Lett.*, 94(23):230504, 2005.

[63] Marcos Curty, Xiongfeng Ma, Bing Qi, and Tobias Moroder. Passive decoy-state quantum key distribution with practical light sources. *Phys. Rev. A*, 81:022310, Feb 2010.

[64] Guifré Vidal. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.*, 91:147902, Oct 2003.

[65] Ian Affleck, Tom Kennedy, Elliott H. Lieb, and Hal Tasaki. Valence bond ground states in isotropic quantum antiferromagnets. *Communications in Mathematical Physics*, 115(3):477–528, 1988.

[66] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac. Matrix product state representations. *Quantum Info. Comput.*, 7(5):401–430, July 2007.

[67] D. Gross, J. Eisert, N. Schuch, and D. Perez-Garcia. Measurement-based quantum computation beyond the one-way model. *Phys. Rev. A*, 76:052315, Nov 2007.

[68] D. Gross and J. Eisert. Quantum computational webs. *Phys. Rev. A*, 82:040303, Oct 2010.

[69] F. D. M. Haldane. Nonlinear field theory of large-spin heisenberg antiferromagnets: Semiclassically quantized solitons of the one-dimensional easy-axis néel state. *Phys. Rev. Lett.*, 50:1153–1156, Apr 1983.

[70] T Kennedy. Exact diagonalisations of open spin-1 chains. *Journal of Physics: Condensed Matter*, 2(26):5737, 1990.

[71] J Ignacio Cirac and Frank Verstraete. Renormalization and tensor product states in spin chains and lattices. *Journal of Physics A: Mathematical and Theoretical*, 42(50):504004, 2009.

[72] B. Zeng S. D. Bartlett J. Lavoie, R. Kaltenbaek and K. J. Resch. Optical one-way quantum computing with a simulated valence-bond solid. *Nature Phys.*, (6):850–854, 2010.

[73] Tomoyuki Morimae and Keisuke Fujii. Not all physical errors can be linear cptp maps in a correlation space. *Scientific Reports*, 2(508).

[74] Lance Fortnow, John Rompel, and Michael Sipser. On the power of multi-prover interactive protocols. In *Theoretical Computer Science*, pages 156–161, 1988.

[75] Hirotada Kobayashi and Keiji Matsumoto. Quantum multi-prover interactive proof systems with limited prior entanglement. *J. Comput. Syst. Sci.*, 66(3):429–450, May 2003.

[76] Keisuke Fujii Tomoyuki Morimae. Blind quantum computation for alice who does only measurements. arXiv:1201.3966, 2012.

[77] Tomoyuki Morimae. How to upload a physical quantum state into correlation space. *Phys. Rev. A*, 83:042337, Apr 2011.

[78] From a personal communication with Scott Aaronson, MIT, USA.

[79] Elham Kashefi Anne Broadbent, Joseph Fitzsimons. $qmip = mip^*$. arXiv:1004.1130v1, 2010.

[80] Dorit Aharonov and Umesh Vazirani. Is Quantum Mechanics Falsifiable? A computational perspective on the foundations of Quantum Mechanics. June 2012.

[81] From a personal communication with Dominique Unruh, University of Tartu, Estonia.

[82] Michele Mosca and Douglas Stebila. Quantum coins. In Aiden A. Bruen and David L. Wehlau, editors, *Error-Correcting Codes, Finite Geometries and Cryptography*, volume 523 of *Contemporary Mathematics*, pages 35–47, Providence, RI, 2010. American Mathematical Society.

[83] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. *Foundations of Computer Science, IEEE Annual Symposium on*, 0:136, 2001.

[84] Dominique Unruh. Universally composable quantum multi-party computation. In *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 486–505. Springer, May 2010. Preprint on arXiv:0910.2912 [quant-ph].

[85] Dominique Unruh. Simulatable security for quantum protocols, September 2004. Preprint on arXiv:quant-ph/0409125.

[86] C. H. Bennett and G. Brassard. Quantum cryptography: Public-key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984.

[87] Tomas Sander, Adam Young, and Moti Yung. Non-interactive cryptocomputing for nc1. In *In 40th Annual Symposium on Foundations of Computer Science*, pages 554–567. IEEE, 1999.

[88] Jean-Sébastien Coron, Avradip Mandal, David Naccache, and Mehdi Tibouchi. Fully homomorphic encryption over the integers with shorter public keys. In *Proceedings of the 31st annual conference on Advances in cryptology*, CRYPTO'11, pages 487–504, Berlin, Heidelberg, 2011. Springer-Verlag.

[89] Craig Gentry. Computing arbitrary functions of encrypted data. *Commun. ACM*, 53(3):97–105, March 2010.

[90] Niel de Beaudrap, Vincent Danos, and Elham Kashefi. Phase map decomposition for unitaries. *(quant-ph/0603266),*, 2006.

[91] N. de Beaudrap, V. Danos, E. Kashefi, and M. Roetteler. Quadratic form expansions for unitaries. In *Theory of Quantum Computation, Communication, and Cryptography Third Workshop, TQC 2008 Tokyo, Japan*, number 5106 in Lecture Notes in Computer Science, 2008.

[92] E. Kashefi, D. Markham, M. Mhalla, and S. Perdrix. Information flow in secret sharing protocols. EPTCS, 9:87, 2009.

[93] M. Van den Nest, W. Dür, G. Vidal, and H. J. Briegel. Classical simulation versus universality in measurement-based quantum computation. *Phys. Rev. A*, 75:012337, Jan 2007.

[94] Jens Eisert and Hans J. Briegel. Schmidt measure as a tool for quantifying multiparticle entanglement. *Phys. Rev. A*, 64:022306, Jul 2001.

[95] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, January 1983. Written in 1969.

[96] Claude E. Shannon. Communication Theory of Secrecy Systems. *Bell Systems Technical Journal*, 28:656–715, 1949.

[97] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography. In *SIAM Journal on Computing*, pages 542–552, 2000.

[98] Mark N. Wegman and J. Lawrence Carter. New classes and applications of hash functions. *Foundations of Computer Science, IEEE Annual Symposium on*, 0:175–182, 1979.

[99] Lars Lydersen, Carlos Wiechers, Christoffer Wittmann, Dominique Elser, Johannes Skaar, and Vadim Makarov. Hacking commercial quantum cryptography systems by tailored bright illumination. *Nature Photonics*, 4(10):686–689, August 2010.

[100] Ilja Gerhardt, Qin Liu, Antia Lamas-Linares, Johannes Skaar, Christian Kurtsiefer, and Vadim Makarov. Full-field implementation of a perfect eavesdropper on a quantum cryptography system. *Nat Commun*, 2:349+, June 2011.

[101] Dominique Unruh. Everlasting quantum security. *IACR Cryptology ePrint Archive*, 2012:177, 2012.

[102] Oded Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.

[103] From a personal communication with Shahram Mossayebi, Royal Holloway University of London.

[104] Michele Mosca, Douglas Stebila, and Berkant Ustaoglu. Quantum Key Distribution in the Classical Authenticated Key Exchange Framework. June 2012.

[105] L. Lamport. Constructing digital signatures from a one-way function. Technical report, October 1979.

[106] P.D. Townsend, J.G. Rarity, and P.R. Tapster. Single photon interference in 10 km long optical fibre interferometer. *Electronics Letters*, 29(7):634 –635, april 1993.

[107] Haruhisa Soda, Ken ichi Iga, Chiyuki Kitahara, and Yasuharu Suematsu. Gainasp/inp surface emitting injection lasers. *Japanese Journal of Applied Physics*, 18(12):2329–2330, 1979.

[108] W. H. Steel. *Interferometry*. Cambridge University Press, 1986.

[109] Patrick J Clarke, Robert J Collins, Philip A Hiskett, Mara-Jos Garca-Martnez, Nils J Krichel, Aongus McCarthy, Michael G Tanner, John A O'Connor, Chandra M Natarajan, Shigehito Miki, Masahide Sasaki, Zhen Wang, Mikio Fujiwara, Ivan Rech, Massimo Ghioni, Angelo Gulinatti, Robert H Hadfield, Paul D Townsend, and Gerald S Buller.

Analysis of detector performance in a gigahertz clock rate quantum key distribution system. *New Journal of Physics*, 13(7):075008, 2011.

[110] A.C. Dada, J. Leach, G.S. Buller, M. Padgett, and E. Andersson. Oexperimental high-dimensional two-photon entanglement and violations of generalized bell inequalities. *Nature Phys.*, 7(9):677–680, 2011.

[111] K.J. Gordon, V. Fernandez, P.D. Townsend, and G.S. Buller. A short wavelength gigahertz clocked fiber-optic quantum key distribution system. *Quantum Electronics, IEEE Journal of*, 40(7):900 – 908, july 2004.

[112] Xu-J. Khan, M. A taxonomy of attacks on quantum key distribution. *International Journal of Latest Trends in Computing*, 2(3), 2011.

[113] C. W. Helstrom. *Quantum detection and estimation theory*. Academic Press, New York, 1976.

[114] R. Renner. Symmetry of large physical systems implies independence of subsystems. *Nature Physics*, 3(9):645–649, 2007. http://arxiv.org/abs/quant-ph/0703069.

[115] M. Christandl, R. Koenig, and R. Renner. Postselection technique for quantum channels with applications to quantum cryptography. *Physical Review Letters*, 102(2):020504, 2009. http://arxiv.org/abs/0809.3019.

[116] Ueli Maurer, Renato Renner, and Stefan Wolf. Unbreakable keys from random noise. In P. Tuyls, B. Skoric, and T. Kevenaar, editors, *Security with Noisy Data*, pages 21–44. Springer-Verlag, 2007.

[117] U.M. Maurer. Secret key agreement by public discussion from common information. *Information Theory, IEEE Transactions on*, 39(3):733 –742, may 1993.

[118] A note on gottesman-chuang quantum signature scheme, 2010. last revised 31 May 2010.

[119] Lawrence M. Ioannou and Michele Mosca. Unconditionally-secure and reusable public-key authentication. August 2011.

[120] Goichiro Hanaoka, Junji Shikata, Yuliang Zheng, and Hideki Imai. Unconditionally secure digital signature schemes admitting transferability. In *In Proc. ASIACRYPT00, Kyoto, December 37*, pages 130–142. Springer-Verlag, 2000.

[121] Goichiro Hanaoka, Junji Shikata, Yuliang Zheng, and Hideki Imai. Efficient and unconditionally secure digital signatures and a security analysis of a multireceiver authentication code. In *Proceedings of the 5th International Workshop on Practice and Theory in Public Key Cryptosystems: Public Key Cryptography*, PKC '02, pages 64–79, London, UK, UK, 2002. Springer-Verlag.

[122] Lu-Ming Duan and Guang-Can Guo. Probabilistic cloning and identification of linearly independent quantum states. *Phys. Rev. Lett.*, 80:4999–5002, Jun 1998.

[123] Franck Ferreyrol, Marco Barbieri, Rémi Blandino, Simon Fossier, Rosa Tualle-Brouri,

and Philippe Grangier. Implementation of a nondeterministic optical noiseless amplifier. *Phys. Rev. Lett.*, 104:123603, Mar 2010.

[124] G. Y. Xiang, T. C. Ralph, A. P. Lund, N. Walk, and G. J. Pryde. Noiseless Linear Amplification and Distillation of Entanglement. *Nature Photonics*, 4:316, 2010.

[125] Fiurasek J. Zavatta A. and Bellini M. A high-fidelity noiseless amplifier for quantum light states. *Nature Photonics*, 5:52, 2011.

[126] Menzies D. and Croke S. Noiseless linear amplification via weak measurements, 2009. arXiv:0903.4181.

[127] John Jeffers. Optical amplifier-powered quantum optical amplification. *Phys. Rev. A*, 83:053818, May 2011.

[128] I. D. Ivanovic. How to differentiate between non-orthogonal states. *Phys. Lett. A*, 123:257, 1987.

[129] D. Dieks. Overlap and distinguishability of quantum states. *Phys. Lett. A*, 126:303, 1988.

[130] A. Peres. How to differentiate between non-orthogonal states. *Phys. Lett. A*, 128:19, 1988.

[131] Anthony Chefles and Stephen M. Barnett. Optimum unambiguous discrimination between linearly independent symmetric states. *Phys. Lett. A*, 250:223, 1998.

[132] Anthony Chefles and Stephen M. Barnett. Quantum state separation, unambiguous discrimination and exact cloning. *J.Pnys.A*, 31:10097, 1998.

[133] Anthony Chefles, Richard Jozsa, and Andreas Winter. On the existence of physical transformations between sets of quantum states, 2003. quant-ph/0307227.

[134] Anthony Chefles. Quantum operations, state transformations and probabilities. *Phys. Rev. A*, 65(5):052314, May 2002.

[135] Yonina C. Eldar, Alexandre Megretski, and George C. Verghese. Optimal detection of symmetric mixed quantum states. *IEEE Transactions on Information Theory*, 50(6):1198–1207, 2004.

[136] Geraldo A. Barbosa, Eric Corndorf, Prem Kumar, and Horace P. Yuen. Secure communication using mesoscopic coherent states. *Phys. Rev. Lett.*, 90:227901, Jun 2003.

[137] Denis Sych and Gerd Leuchs. Coherent state quantum key distribution with multi letter phase-shift keying. *New Journal of Physics*, 12(5):053019, 2010.

[138] *Matrix Computations (Johns Hopkins Studies in Mathematical Sciences)(3rd Edition)*. The Johns Hopkins University Press, 3rd edition, 1996.

[139] David T. Pegg, Lee S. Phillips, and Stephen M. Barnett. Optical state truncation by projection synthesis. *Phys. Rev. Lett.*, 81(8):1604–1606, Aug 1998.

[140] S. A. Babichev, J. Ries, and A. I. Lvovsky. Quantum scissors: teleportation of single-mode optical states by means of a nonlocal single photon. *Europhys. Lett.*, 64:1, 2003.

[141] Mario A. Usuga, Christian R. Muller, Christoffer Wittmann, Petr Marek, Radim Filip, Christoph Marquardt, Gerd Leuchs, and Ulrik L. Andersen. Noise-powered probabilistic concentration of phase information. *Nature Physics*, 6(10):767–771, August 2010.

[142] P. Marek R. Filip C. Marquardt G. Leuchs C. R. Müller, C. Wittman and U. L. Andersen. Probabilistic cloning of coherent states without a phase reference, 2011. arXiv:1108.4241.

[143] A. Chefles. How to differentiate between non-orthogonal states. *Phys. Lett. A*, 239:339, 1998.

[144] B. Huttner, N. Imoto, N. Gisin, and T. Mor. Quantum cryptography with coherent states. *Phys. Rev. A*, 51:1863–1869, Mar 1995.

[145] S. J. van Enk. Unambiguous state discrimination of coherent states with linear optics: Application to quantum cryptography. *Phys. Rev. A*, 66:042313, Oct 2002.

[146] Stephen M. Barnett and Sarah Croke. Quantum state discrimination. *Adv. Opt. Photon.*, 1(2):238–278, Apr 2009.

[147] János A. Bergou. Tutorial review. Discrimination of quantum states. *J. Mod. Opt.*, 57(3):160–180, 2010.

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

[149] V. Bužek and M. Hillery. Quantum copying: Beyond the no-cloning theorem. *Phys. Rev. A*, 54:1844–1852, Sep 1996.

[150] A. Peres. How the no-cloning theorem got its name. 2002.

[151] Log-concave probability and its applications, 1989.

[152] Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26:445–469, 2005. 10.1007/s00199-004-0514-4.

[153] P. Wocjan and S. Zhang. Several natural BQP-complete problems. 2006.