UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Resource requirements for quantum cryptography

Muguruza Lasa, G.

Publication date
2025

Document Version
Final published version

Link to publication

Citation for published version (APA):
Muguruza Lasa, G. (2025). *Resource requirements for quantum cryptography*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:14 Oct 2025

In this thesis we explore communication between parties with access to quantum resources, such as channels, qudits, and computers.

We start by studying types of quantum channels. In particular, we consider a scenario where the sender knows a classical description of the qudit they intend to send, and the receiver's operations are restricted to classical ones. Our main result is that the accuracy of the transmission scales inverse exponentially with the number of pre-shared entangled qudits.

We later look into possible extra properties of quantum channels by giving a protocol for authenticating a noisy channel. Moreover, we prove that our protocol requires access to poly-logarithmic fewer qubits than the previously known techniques.

For the rest of the dissertation we look at what ideal quantum channels could be useful for. Our first result is a round-optimal quantum protocol for oblivious transfer which can be instantiated both in the plain and quantum random oracle models (by basically lifting the properties of an underlying zero-knowledge protocol), but we obtain round optimality in the quantum random oracle model only.

The last chapters of this dissertation are an exploration of a particular set-up assumption, quantum pseudorandomness; i.e. assuming the existence of random (looking) quantum states. We first show that in contrast to the classical case, the size of a quantum pseudorandom object cannot be shrunk. Finally, we prove that if there is a promise problem that admits a quantum reduction that loses information about its input, then certain quantum pseudorandom primitives exist.

Garazi Muguruza Lasa (1997) conducted her research at QuSoft and MNS, within CWI and the University of Amsterdam. She previously completed a master's at TU Delft and a bacheor's at Universidad Complutense de Madrid. Born in Azpeitia (Euskal Herria), she currently lives in Delft (Netherlands) where she enjoys running in nature and collaging with friends.

Resource requirements for Quantum Cryptography

Garazi Muguruza Lasa

Resource requirements for Quantum Cryptography

Garazi Muguruza Lasa

Resource requirements for Quantum Cryptography

Garazi Muguruza Lasa

Copyright © by Garazi Muguruza Lasa.

Cover design by Garazi Muguruza Lasa. It is a collage of the mountain Anboto above Berain, with the cave of Mari in between.

Resource requirements for Quantum Cryptography

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. P.P.C.C. Verbeek
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Aula der Universiteit
op vrijdag 17 oktober 2025, te 11.00 uur

door Garazi Muguruza Lasa
geboren te Azpeitia

Promotiecommissie

Promotores:	prof. dr. P. Grosso	Universiteit van Amsterdam
	prof. dr. C. Schaffner	Universiteit van Amsterdam
Copromotores:	dr. F. Speelman	Universiteit van Amsterdam
Overige leden:	prof. dr. A. Broadbent	University of Ottawa
	dr. A. Bredariol Grilo	Sorbonne University-CNRS
	prof. dr. S.M. Jeffery	Universiteit van Amsterdam
	dr. N.A. Resch	Universiteit van Amsterdam
	prof. dr. C.J.M. Schoutens	Universiteit van Amsterdam

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Contents

List of publications

The content of this dissertation is based on the following articles. The author names are ordered alphabetically, and all authors contributed equally.

[CMS23] **Oblivious transfer from zero-knowledge proofs.** Léo Colisson Palais, Garazi Muguruza and Florian Speelman. In *Advances in cryptology—ASIACRYPT 2023.*

[DMS25] **An efficient combination of quantum error correction and authentication.** Yfke Dulek, Garazi Muguruza and Florian Speelman. In *IACR Communications in Cryptology 1 (4).*

[MS24] **Port-Based State Preparation and Applications.** Garazi Muguruza and Florian Speelman. In *Quantum 8, 1573.*

[BEM24] **Quantum Pseudorandomness Cannot Be Shrunk In a Black-Box Way.** Samuel Bouaziz–Ermann and Garazi Muguruza. *ArXiv preprint.*

[FGMR25] **Cryptography from Lossy Reductions.** Pouria Fallahpour, Alex B. Grilo, Garazi Muguruza and Mahshid Riahinia. *ePrint preprint.*

During the course of her PhD, the author has additionally co-authored the following work, which is not included in this dissertation:

[EGM24] **All S_p notions of quantum expansion are equivalent.** Francisco Escudero Gutiérrez and Garazi Muguruza. *ArXiv preprint.*

Chapter 1

Introduction

Life constantly deploys communication as a means for cooperation; roosters have a special sound to tell hens about worms in the area and pea plants use their root network to inform about possible draughts, to name a couple. At the same time humans are specially busy developing new technologies to increase their communication capacity; from writing systems to the internet.

One of the most recent technologies, at the time of writing this thesis, are quantum technologies. Once we agreed that the world probably works following the rules of quantum mechanics, we are busying ourselves trying to understand its consequences. What is already clear is that we require a complete shift in how we imagine communication.

To be able to effectively cooperate between scientists from various backgrounds, we need to start from finding a common language. For each letter I write in this thesis, mostly in English, I have to choose one amongst 26 possible ones. Actually, each time, for each possible letter, I have to decide to either use it or not, and only stop when I choose to use one. This amounts for a total of 26 decisions per letter. When modelling the minimum information package as a bit (0 for absence and 1 for presence), we can represent each letter by 5 bits. Therefore, with each choice of a letter I can communicate to you up to nearly 5 bits of information. As mentioned earlier, quantum mechanics have nevertheless taught us that bits are not the minimum information package present in Nature, these are instead qubits (which can be in a coherent superposition of all computational states simultaneously).

We have already established what our letters will be, qubits, but we now need a way to send them around. One fascinating phenomena that follows the new theory of information is teleportation; a technique for transferring quantum information using purely classical channels and statistical correlations. For teleportation to take place we need the sender and the receiver to share an entangled state, two maximally correlated qubits, and the sender to hold an extra qubit of her own. Note that two spatially separated quantum states can still be arbitrarily correlated. The sender will then perform a measurement on both her qubits, at which point the receiver's state will immediately change due to the shared correlations, but the receiver will not be able to obtain the correct state until the sender communicates two extra bits of information.

Established that we can probably, given enough entangled pairs, share quantum information with each other, more human aspects of communication enter the picture. For this technology to be useful, we would want to be able to choose who receives our messages, and who we want to keep them hidden from. We develop social relations by manipulating ratios of knowledge and ignorance; for example, football players and coaches at Athletic speaking Euskara, a language isolate[1] from Euskal Herria spoken by barely a million inhabitants, are able to hide their tactics from the opponents. Decrypting an existing language is relatively easy though, the opponents need only to learn it in its spoken form.

Written language offers many more avenues for secrecy, and thus gives birth to the art of keeping information private, Cryptography. The first known protocol for written secrecy dates from Roman times. Based on the principle of every word being a choice of letters, we can shift one to the right a letter in the alphabet so that words such as *kfmmp* have no meaning until deciphered to *hello*. The shifting to the right technique, or the Caesar cipher, is a protocol and the exact number of shifts, one in this case, is the key of the protocol. Once the number of shifts and the language of the original texts are known, discovering the correct key to decrypt the Caesar cipher only takes some pattern matching effort. Not only that, once someone finds which shift was used for encrypting the information, they are able to read the whole message sent, regardless of its length.

Defining what it means for a message to be truly secret is far from straightforward. Intuitively, we would like a person without access to the key to learn absolutely nothing about the original message, except perhaps its length. But

[1] Language isolates refer to languages that are not related to any other language or language family. You can find more information about them in [SKSZ25].

this stringent notion of perfect secrecy comes at a steep price, as it forces encryption keys to be as long as the messages themselves. To escape this limitation, cryptography takes a pragmatic turn: rather than demanding that ciphertexts reveal zero information about the original message, we instead require that extracting any meaningful insight from them remains *very hard*.

Naturally, this leads us to the next hurdle: what exactly makes a problem *hard*? Take learning the Marathi language, for example; it seems hard at first glance, with an alphabet more than twice the size of the English language's, yet over a hundred million people speak it fluently, proving it is ultimately possible. In the context of computer science, we say that a problem is hard if no computer can solve it within practical time limits, and easy if it can. For cryptography, not only we need problems that are inherently hard, we need ones that become easy for anyone holding the right key.

But, are there any hard problems in Nature? Cryptographers come up with various hypotheses about the status of Nature, famously Impagliazzo's five classical worlds [Imp95] and an extra quantum one [BMM$^+$25].

i. **Cryptomania:** the ideal world for cryptography, where parties can securely communicate without previously having to agree on a secret key.

ii. **Minicrypt:** a world where hard-on-average but easy given the right key problems exist, however, parties cannot agree on a secret using only public channels.

iii. **Algorithmica, Heuristica and Pessiland:** worlds without classically hard-on-average problems, and hence not much interesting cryptography.

iv. **Microcrypt:** the first purely quantum world, where random-looking quantum states can be efficiently sampled, and hence quantum cryptography is possible, even if classically hard-on-average problems do not exist.

Proving in which world we live is a problem that will probably require much cooperation, and communication.

Quantum Communication. We mentioned earlier how communicating quantum information between two quantum computers connected just via classical channels is possible, as long as they pre-share entanglement. Then again, statistical correlations between qubits are not their only intriguing property. Another fascinating quality of quantum information is that even though you might hold a qubit in an arbitrary superposition of two states, not all this information is *accessible* to you. To know if you write the number 0 or 1, at the end of the day, you will need to look at your quantum state, i.e. *measure* it, and although until you looked the state was simultaneously both numbers 0 and 1, the moment you look, you will only see either one of them. In the original quantum teleportation protocol, the sender, we will call her *Anboto* from now on, does not necessarily know the quantum state she holds, she has not looked at it yet. There is another variant of teleportation, called *Remote State Preparation* (RSP), where Anboto knows the qubit she intends to send, that is, she can give a classical description of it. Naturally, if she knows the state she could also prepare it and just use the original teleportation protocol to transmit it, and indeed since Anboto has more information in RSP, it is easier to perform. Specifically, she only requires to send 1 bit of information per qubit asymptotically, 1 bit less per qubit than in the original teleportation protocol.

A different approach to studying communication is to find intermediate scenarios between parties with classical and quantum capabilities. For example, even being able to keep a qubit stable in the lab is currently a non-trivial task, not to mention operating on these. *Port-Based Teleportation* (PBT) is a variant of teleportation where Anboto is trying to communicate with *Berain*; who does not have quantum computational power and can only store quantum states. We still assume that both parties pre-share arbitrary entangled qubits. In such scenario, it is not possible to *always perfectly* share a qubit! And therefore we need to consider what is the goal of communicating: does Anboto want to *always* communicate a qubit, even if it is not exactly the one she intended but *very similar* to it, or does she instead want only succeed *sometimes* but have the guarantee that when she succeeds Berain will hold exactly the intended state?

This difficulty in defining what exactly the goal of our tasks should be, and how to measure success, will follow us throughout the whole scientific endeavour.

Quantum Authentication. Once the quantum communication channel is established, we need to make sure that it is secure. One of the properties we can ask of a channel is *authentication*: making sure that the encrypted message has not been changed during the process. The channel will naturally have some noise, so small changes are unavoidable, but how do we distinguish if the changes are due to just noise or an active intermediary, who intercepted and replaced the information we sent?

Turns out that authenticating quantum data is at least as hard as encrypting it [BCG$^+$02], but it is possible to do so without assuming the hardness of any problem, as long as the honest participants have access to a pre-shared symmetric key.

Quantum Oblivious Transfer. Quantum cryptography is not only useful for securely communicating quantum information, but we can also consider quantum protocols for ensuring security of classical primitives, even in the presence of quantum eavesdroppers.

For example, imagine you are trying to set up a protocol to ensure safe workplace interactions. You want to provide as much information as possible to the employees, but they are afraid that if they make multiple queries, e.g. think of "sexism" followed by "xenophobia", their intersectional identity might be traced. Therefore, you want to ensure that when they make a query, they receive the corresponding information without revealing what they asked. This seemingly simple task is called *oblivious transfer*: a cryptographic primitive in which a sender transfers only one of potentially many secrets to a receiver, while remaining oblivious as to which piece was chosen by the receiver.

There are multiple ways of obtaining quantum-secure oblivious transfer; but they all require *hardness* assumptions, that is, restrictions on the adversary.

Quantum Pseudorandomness. The last potential state of Nature, Microcrypt, emerged as a research field at the same period as my doctoral work, and hence is specially dear to me.

On the one hand, since we already agreed that our existence follows the laws of quantum mechanics, we now wonder if secrecy can be a consequence of *quantum hardness* only. We can define quantum hardness in many ways; the hardness of distinguishing between a pseudorandom state and a random state, the hardness of inverting a one-way state generator, the hardness of distinguishing two efficiently samplable states that are far from each other in trace distance, etc. Although the classical counterparts of all these primitives are equivalent, we lack provable connections between the quantum ones. Even more, we have evidence that suggests many of the primitives are *separated* from each other, meaning that one cannot build one from the other.

Remarkably, quantum pseudorandom primitives can exist even in the absence of classically hard-to-verify problems, yet we still require problems that are quantumly hard-to-solve but easy-to-verify. The difference between solving and verifying a problem appears intuitive; it is very hard to construct a rigorous mathematical proof, but checking the solution's validity is comparatively easy. It does then perhaps not come as a surprise that coming up with problems that are quantumly hard-to-solve but easy-to-verify, is indeed very difficult. Canonical examples, like finding the prime factorization of an integer, unfortunately seem easy for Nature!

Instead of directly coming up with such examples, we can try to define properties of a problem that would make them hard.

1.1 Our contributions

This thesis is divided in six chapters, including a preliminary chapter and five content chapters, each exploring a different aspect of secure communication in a quantum world.

Chapter 2: Preliminaries. In this background chapter we introduce the common notation for the rest of the thesis. The chapter includes a background section that introduces the two security frameworks we will consider in Chapters 4 and 5. The concepts are not original but a concise presentation notions taken from [HSS11] and [Por17].

Chapter 3: Port-Based State Preparation. The third chapter of this thesis is based on joint work with Florian Speelman [MS24].

In this chapter we introduce Port-Based State Preparation (PBSP), a teleportation task where Anboto holds a complete classical description of the target state and Berain's correction operations are restricted to only tracing out registers. We show a protocol that implements PBSP with error decreasing exponentially in the number of ports, in contrast to the polynomial trade-off for the related task of Port-Based Teleportation, and we prove that this is optimal when a maximally entangled resource state is used.

As an application, we introduce approximate Universal Programmable Hybrid Processors (UPHP). Here the goal is to encode a unitary as a quantum state, and the UPHP can apply this unitary to a quantum state when knowing its classical description. We give a construction that needs strictly less memory in terms of dimension than the optimal approximate Universal Programmable Quantum Processor achieving the same error. Additionally, we provide lower bounds for the optimal trade-off between memory and error of UPHPs.

Compared to previous works on related topics, our proofs are very simple and rely on basics of information theory such as the no-go theorem for quantum random access codes and the no-signalling principle.

Chapter 4: The Threshold Authentication Code. The fourth chapter of this thesis is based on joint work with Yfke Dulek and Florian Speelman [DMS25].

When sending quantum information over a channel, we want to ensure that the message remains intact. Quantum error correction and quantum authentication both aim to protect (quantum) information, but approach this task from two very different directions: error-correcting codes protect against probabilistic channel noise and are meant to be very robust against small errors, while authentication codes prevent adversarial attacks and are designed to be very sensitive against any error, including small ones.

In practice, when sending an authenticated state over a noisy channel, one would have to wrap it in an error-correcting code to counterbalance the sensitivity of the underlying authentication scheme. We study the question of whether this can be done more efficiently by combining the two functionalities in a single code. To illustrate the potential of such a combination, we design the threshold code, a modification of the trap authentication code which preserves that code's authentication properties, but which is naturally robust against depolarizing channel noise. We show that the threshold code needs polylogarithmically fewer qubits to achieve the same level of security and robustness, compared to the naive composition of the trap code with any concatenated CSS code.

We prove the security of our code in the abstract cryptography framework, which ensures the universal composability of our code without any hardness assumptions on the participants, as long as they pre-share a secret key.

Chapter 5: Round-Optimal Oblivious Transfer. The fifth chapter of this thesis is based on joint work with Léo Colisson Palais and Florian Speelman [CMS23]. Although there are more than the mentioned results in the original work, the following subset reflects my contributions.

We provide a generic construction to turn any classical zero-knowledge (ZK) protocol into a composable quantum oblivious transfer (OT) protocol, mostly lifting the round-complexity properties and security guarantees of the underlying ZK protocol to the resulting OT protocol. In particular, by instantiating our construction using non-interactive ZK, we provide the first round-optimal (2-message) quantum OT protocol secure in the random oracle model.

At the heart of our construction lies a way of proving that a received quantum state is in a particular basis without revealing additional information on it, even in a non-interactive way, without public-key primitives, and/or with statistical guarantees when using an appropriate classical ZK protocol.

Chapter 6: Background on quantum Pseudorandomness. The sixth chapter of this thesis is an original introduction to quantum pseudorandomness, not published elsewhere.

This chapter provides an overview of the current state of the field and examines the key properties of cryptographic primitives that are essential for understanding the significance of the final two chapter of the thesis, namely Chapters 7 and 8.

Chapter 7: No shrinking of Quantum Pseudorandomness. The seventh chapter of this thesis is based on joint work with Samuel Bouaziz–Ermann [BEM24].

In this chapter we show that there exists a quantum oracle relative to which pseudorandom state generators (PRSGs) exist but PRSGs with log-length output (short-PRSGs) do not. The oracle of the separation is a collection of random unitaries introduced by Kretschmer [Kre21], and insights into pseudo-determinism.

Chapter 8: Construction of Quantum Pseudorandomness. The last chapter of this thesis is based on join work with Pouria Fallahpour, Alex B. Grilo and Mahshid Riahinia [FGMR25]. Although there are more than the mentioned results in the original work, the following subset reflects my contributions.

In this chapter we study complexity-theoretical hardness conditions for the existence of quantum pseudorandomness. In particular, we explore the relation between quantum pseudorandom primitives and lossy reductions, i.e., reductions R from classical bit-strings to quantum states, for which it holds that the $I_q(X; R(X)) \ll \lambda$ for all distributions X over inputs of size λ. We give two results in this direction; firstly, we prove that either efficiently generatable statistically far but computationally indistinguishable pairs of states

(EFIs) exist or the run-time of a worst-case to average case f-distinguisher reductions for a promise problem Π must be *high*. Secondly, we show an analogous result for one-way state generators (OWSGs).

We do this by showing that the aforementioned reductions are a generalized version of lossy reductions, lossy just for *sparse uniform* distributions (which we call mildly-lossy). The main insight is an extension of the *disguising distribution lemma* by Drucker [Dru12] to include mildly-lossy reductions. We also show that the results extend to f-reductions as long as f is a non-constant permutation-invariant Boolean function, which includes And-, Or-, Maj-, Parity-, Mod_k- and Threshold_k-reductions.

Chapter 2

Preliminaries

In this chapter we include key background information to contextualize the rest of the thesis.

In Section 2.1 we establish the notational conventions used throughout the rest of the work, which serves as a centralized reference for later chapters. In Section 2.2 we introduce the two security frameworks we will be working on: the stand-alone security model (employed in Chapter 5) and the abstract cryptography model (used in Chapter 4). While the core content of these sections derives from articles [CMS23] and [DMS25] respectively, we include new comparative insights to highlight their differences.

2.1 Notation

As is customary in multi-party communication systems, we give real names to participants for convenience and to aid comprehension. I have chosen to name them *Anboto* and *Berain* to reflect my love for the present work; Anboto is the name of a mountain in Bizkaia where the personification of mother nature *Mari* lives, and Berain is the name of another mountain in Nafarroa, at the heart of Euskal Herria. Both are common Basque names.

For a positive integer n, we let $[n]$ denote the set $\{1, 2, \ldots, n\}$. The set of natural numbers is denoted by $\mathbb{N}$ and we denote by $\mathbb{R}^+$ the set of positive real numbers. We will use log to denote the base-2 logarithm. The binary entropy function is denoted by h, that is $h(p) := -p\log(p) - (1-p)\log(1-p)$ for any $p \in [0,1]$. Let f and g be non-negative real-valued functions. We say that $f(x) = O(g(x))$ if $|f|$ is asymptotically bounded from above by g, that is, there exists a constant $c > 0$ and a real number x_0 such that

$$|f(x)| \leq c|g(x)| \quad \text{for all } x \geq x_0 \,. \tag{2.1}$$

We say that $f(x) = \omega(g(x))$ if instead g is asymptotically bounded from above by f, that is, for every constant $c > 0$ and there exists a real number x_0 such that

$$f(x) \geq cg(x) \quad \text{for all } x \geq x_0 \,. \tag{2.2}$$

We say that $f(x) = \Omega(g(x))$ if g is asymptotically dominated by f, that is, there exists a constant $c > 0$ and a real number x_0 such that

$$f(x) > cg(x) \quad \text{for all } x \geq x_0 \,. \tag{2.3}$$

We say that $f(x) = \Theta(g(x))$ if f is asymptotically bounded both from above and below by g, that is, there exist constants $c_1, c_2 > 0$ and a real number x_0 such that

$$c_2 g(x) \geq f(x) \geq c_1 g(x) \quad \text{for all } x \geq x_0 \,. \tag{2.4}$$

Probability theory. For a distribution $\mathcal{D}$, $x \leftarrow \mathcal{D}$ denotes that x is sampled according to $\mathcal{D}$. For any finite set S, we let $\mathcal{U}_S$ denote the uniform distribution over the set S. When it is clear from context we might write $x \leftarrow S$ to denote that x is sampled uniformly at random from S.

Lemma 2.1.1 (Multiplicative Chernoff bound)**.** *Let $X_1, \ldots, X_n$ be independent $\{0,1\}$-valued random variables. Let X be their sum and let $\mu = \mathbb{E}[X]$.*

It holds that for any $\delta \in (0,1)$:

$$\Pr(X \leq (1-\delta)\mu) \leq \exp\left(-\frac{\mu\delta^2}{2}\right).$$

Lemma 2.1.2 (Markov's inequality)**.** *Let X be a real positive-valued random variable, and let $a > 0$. Then*

$$\Pr(X \geq a) \leq \mathbb{E}[X]/a\,.$$

Cryptography. A function $f \colon \mathbb{N} \to \mathbb{R}$ is said to be *negligible* if for any positive integer $c \in \mathbb{N}$ there exists $n_0 \in \mathbb{N}$ such that for all $n > n_0$, $|f(n)| < n^{-c}$. We let negl denote an unspecified negligible function, which is by definition (eventually) smaller than any inverse polynomial. We let poly denote an unspecified polynomial. We let $\lambda \in \mathbb{N}$ denote the security parameter, we explain more in Section 2.2.

Definition 2.1.3 (Indistinguishable random variables)**.** *Two sequences of real-valued random variables $\mathbf{X} = \{X_\lambda\}_{\lambda\in\mathbb{N}}$ and $\mathbf{Y} = \{Y_\lambda\}_{\lambda\in\mathbb{N}}$ are said to be* ε-indistinguishable, *denoted $\mathbf{X} \approx_\varepsilon \mathbf{Y}$, if $|\Pr(X_\lambda = 1) - \Pr(Y_\lambda = 1)| \leq \varepsilon(\lambda)$. In particular, if $\varepsilon = \mathsf{negl}(\lambda)$, $\mathbf{X}$ and $\mathbf{Y}$ are said to be indistinguishable, denoted $\mathbf{X} \approx \mathbf{Y}$.*

Computer science. The set $\{0,1\}^*$ contains all bit strings of arbitrary length, that is, $\{0,1\}^* := \cup_{n=1}^{\infty}\{0,1\}^n$. For a bit string $x \in \{0,1\}^*$, we denote its bit-length by $|x|$. We denote its i-th bit starting from 0 by x_i and its i-th bit starting from 1 by $x[i]$. We use $||$ to denote the concatenation operator. For two bits $x, y \in \{0,1\}$, we denote their xor by $x \oplus y = x \cdot y \mod 2$. For two bit strings of the same length $x, y \in \{0,1\}^n$, we denote their inner product by $\langle x, y\rangle := \sum_{i=0}^{n} x_i \oplus y_i$. Throughout this work we assume that all functions used to represent the lengths of the cryptographic primitives are quantum polynomial-time computable. A *language* $\mathcal{L}$ is a subset of $\{0,1\}^*$ and its complement is defined as $\overline{\mathcal{L}} := \{0,1\}^* \setminus \mathcal{L}$. A *promise problem* Π consists of two disjoint sets $\Pi_\mathrm{Y}, \Pi_\mathrm{N} \subseteq \{0,1\}^*$ respectively referred to as the set of YES and NO instances. Problem Π asks to decide whether given an instance $x \in \{0,1\}^*$, which is promised to lie in $\Pi_\mathrm{Y} \cup \Pi_\mathrm{N}$, belongs to Π_Y or Π_N.

Definition 2.1.4 (Characteristic Function of a Promise Problem)**.** *For a promise problem Π, the characteristic function of Π is the map $\chi_\Pi(x)$:*

$\{0,1\}^* \to \{0,1,\star\}$ *given by*

$$\chi_\Pi(x) = \begin{cases} 1 & \textit{if } x \in \Pi_{\mathrm{Y}} \\ 0 & \textit{if } x \in \Pi_{\mathrm{N}} \\ \star & \textit{otherwise} \end{cases} .$$

Quantum information. Given a Hilbert space $\mathcal{H}$ with associated inner product $\langle \cdot, \cdot \rangle$, we denote by *kets* $|\phi\rangle$ the normalized vectors in $\mathcal{H}$ and by *bras* $\langle \phi|$ the corresponding dual vectors in $\mathcal{H}^*$, that is, $\langle \phi| := \langle \cdot|\phi\rangle$. We denote by $\|\cdot\|_{\mathcal{H}}$ the norm defined by the inner product of the Hilbert space $\mathcal{H}$, that is $\|\cdot\|_{\mathcal{H}} := \sqrt{\langle \cdot, \cdot \rangle}$. We define the 1-norm of any Hermitian matrix A on a Hilbert space $\mathcal{H}_d$ as $\|A\|_1 := \mathrm{Tr}\left(\sqrt{A^\dagger A}\right) = \sum_i |\lambda_i|$, where λ_i's are the eigenvalues of A. For a positive, semi-definite matrix A, we have $\|A\|_1 = \mathrm{Tr}(A)$, and therefore $\mathrm{Tr}\left(A^2\right) \leq \mathrm{Tr}(A)^2$. The *Frobenius norm* of an $n \times m$ matrix A with elements $(a_{ij})_{i\in[n],j\in[m]}$ is defined as $\|A\|_F := \sqrt{\sum_i \sum_j |a_{ij}|^2}$. The *trace distance* of two Hermitian matrices A and B refers to the natural associated distance, that is, $\mathrm{TD}(A,B) := \frac{1}{2}\|A-B\|_1$. We use $\mathrm{D}(\mathcal{H}_d)$ to denote (possibly mixed) states in a d-level system, called *qudits*, that is, trace-one positive semidefinite matrices in a d-dimensional Hilbert space

$$\mathrm{D}(\mathcal{H}_d) := \{\rho \in \mathrm{PSD}(\mathcal{H}_d) \colon \ \mathrm{Tr}(\rho) = 1\} . \tag{2.5}$$

A positive operator valued measurement (POVM) is a function μ from a finite set Ω to positive semidefinite matrices over a Hilbert space $\mathcal{H}$ such that

$$\sum_{x\in\Omega} \mu(x) = I , \tag{2.6}$$

it is customary to denote the POVM by $\mathcal{M} = \{M_x\}_{x\in\Omega}$, where $M_x := \mu(x)$. If we measure a state $\rho \in \mathrm{D}(\mathcal{H})$ using measurement μ, the probability of outcome $x \in \Omega$ is given by $\mathrm{Tr}(M_x \rho)$. The trace distance gives the maximum distinguishability of two quantum states, that is,

$$\mathrm{TD}(\rho, \sigma) = \max_{0\leq M\leq I} \mathrm{Tr}(M(\rho - \sigma)) . \tag{2.7}$$

A quantum state is called *pure* if it is of the form $\rho = |\phi\rangle\langle\phi|$ for some unit vector $|\phi\rangle \in \mathcal{H}_d$, in such cases we might denote the state ρ simply by $|\phi\rangle$. We write $\mathbb{S}(d)$ to denote the set of d-dimensional pure quantum states. Two-level systems are called *qubits*. We let MS_n denote the set of all states over n qubits and we define $\mathsf{MS}_* := \cup_{n=1}^{\infty}\mathsf{MS}_n$. We sometimes consider unnormalized quantum states, e.g. $\Pi\rho\Pi$ for some projection Π and quantum state ρ. For tensor product spaces we sometimes use subscripts to indicate which subsystem the mathematical object refers to. For example, $\mathcal{H}_{AB} := \mathcal{H}_A \otimes \mathcal{H}_B$ and we refer to a state in $\mathrm{D}(\mathcal{H}_{AB})$ as ρ_{AB}. We sometimes write registers with bold letters, i.e. $\mathbf{A}, \mathbf{B}$. For a quantum state ρ_{AB} on two subsystems A and B, we denote by ρ_A the state on subsystem A alone, that is, $\rho_A := \mathrm{Tr}_B(\rho_{AB})$, where $\mathrm{Tr}_B \colon \mathrm{L}(\mathcal{H}_A \otimes \mathcal{H}_B) \to \mathrm{L}(\mathcal{H}_A)$ denotes the partial trace operator over subsystem B. For a basis $\{|b\rangle\}_{b\in\mathcal{K}}$ of the Hilbert space $\mathcal{H}_B$, the partial trace acts on ρ_{AB} as

$$\mathrm{Tr}_B(\rho_{AB}) := \sum_{b\in\mathcal{K}} (I_A \otimes |b\rangle_B)\rho_{AB}(I_A \otimes \langle b|_B)\,. \tag{2.8}$$

A *quantum channel* is a linear map $\Phi_{A\to B} \colon \mathrm{L}(\mathcal{H}_A) \to \mathrm{L}(\mathcal{H}_B)$ between quantum states that is completely positive and trace preserving. We denote the set of quantum channels from $\mathcal{H}_A$ to $\mathcal{H}_B$ by $\mathrm{CPTP}(\mathcal{H}_A, \mathcal{H}_B)$. If the input and output Hilbert space $\mathcal{H}$ of a quantum channel are the same, then we instead write $\mathrm{CPTP}(\mathcal{H})$. A quantum channel $\Phi_{A\to B}$ can also be represented by its *Kraus operators*, that is operators $B_1, \dots, B_k \in \mathrm{L}(\mathcal{H}_A, \mathcal{H}_B)$ such that $\sum_{i=1}^k B_i^\dagger B_i = I_A$ and

$$\Phi_{A\to B}(M) = \sum_{i=1}^{k} B_i M B_i^\dagger \quad \text{for all} \quad M \in \mathrm{L}(\mathcal{H}_A)\,. \tag{2.9}$$

For an operator $\Phi_{A\to B} \in \mathrm{L}(\mathcal{H}_A, \mathcal{H}_B)$, we let $\|\Phi_{A\to B}\|_{op}$ denote its operator norm, that is, $\|\Phi_{A\to B}\|_{op} := \max_{x\in\mathcal{H}_A} \|\Phi(x)\|_{\mathcal{H}_B}$. We drop the subscript when the spaces the operator is acting on are clear from context. The single-use distinguishability between two quantum channels is often quantified by the distance induced by the diamond norm, given $\Phi \in \mathrm{L}(\mathcal{H}_A, \mathcal{H}_B)$, we denote $\|\Phi\|_\diamond := \sup_{k\in\mathbb{N}} \|I_{\mathbb{C}^k} \otimes \Phi\|_{op}$. The diamond-norm is subadditive under composition, that is, for two quantum maps $\Phi \in \mathrm{L}(\mathcal{H}_A, \mathcal{H}_B)$ and $\Psi \in \mathrm{L}(\mathcal{H}_B, \mathcal{H}_C)$ we have $\|\Psi \circ \Phi\|_\diamond \le \|\Psi\|_\diamond \|\Phi\|_\diamond$. A quantum channel Φ is in particular a bounded operator, thus does not increase the trace norm, that is for any Hermitian matrix A, it holds that $\|\Phi(A)\|_1 \le \|A\|_1$. Therefore, for any state ρ (even unnormalized), we have $\mathrm{Tr}(\Phi(\rho)) \le \mathrm{Tr}(\rho)$.

Quantum computing. Let us denote the set of unitary operations in a d-dimensional Hilbert space by $\mathrm{U}(d)$. A particularly interesting subset of quantum channels for quantum computing are *unitary transformations*, consisting of a single Kraus operator $U \in \mathrm{U}(d)$. The following *Pauli matrices* form a basis of single-qubit Hermitian operations,

$$I := \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, X := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, Z := \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, Y := \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} = iXZ\,. \tag{2.10}$$

Note that any two Pauli matrices either commute or anti-commute.

2.2 Semantic security in a Quantum World

Modern cryptography requires tasks to be precisely defined in the mathematical sense, in order to show that they are *provably secure*, that is, we can analyse them rigorously relative to clearly stated assumptions. At their core, most cryptographic schemes consist of efficient algorithms —such as key generation, encryption and verification— that must satisfy both *correctness* and *security* guarantees. Given a *security parameter* $\lambda \in \mathbb{N}$, we say that an algorithm is efficient if its running time is a polynomial function in λ, which we give as a unary argument.

Correctness requires that the protocol fulfils the task as intended for honest users. For example, in an encryption scheme, decrypting an encrypted message should recover the original plaintext. While some schemes achieve perfect correctness, others may allow negligible error probability in the security parameter.

Representing security is a more subtle task. Both characterizations of security we will use in this work are within the *simulation* paradigm involving a *real world* and an *ideal world.* We characterize the desired task as an interactive machine $\mathcal{F}$ called a *functionality*, which can be either classical or quantum. The real world represents an actual run of the protocol Π where some parties can potentially be corrupted, while the ideal world represents an idealized version of the protocol where the parties are only allowed to interact through the trusted ideal functionality. For the protocol to be deemed secure, anything that an adversary can learn or do in an interaction with the actual protocol (in the real world) should be simulatable in an interaction with this ideal functionality (in the ideal world). The environment attempts to distinguish the real and ideal worlds, and should not be able to tell the difference between the two worlds with non-negligible probability. If both worlds are indistinguishable, then the protocol is said to be secure as any attack doable in the real world would apply in the ideal world (otherwise it would provide a way to distinguish both worlds) and therefore to the ideal functionality, which is secure by definition. A negligible probability of success in distinguishing the two worlds is considered a safe margin because even polynomially many repetitions made by an efficient adversary would not increase the probability of success to a non-negligible distinguishability.

Nevertheless, in the following chapters we will work with two different security models, that of quantum stand-alone security introduced by Hallgren, Smith and Song [HSS11], and the Abstract Cryptography framework by Maurer and Renner [MR11]. In the former model, security is built bottom-up, starting from computational models and describing how the machines can interact with each other. Moreover, the stand-alone framework provides a proof of security when the protocols are run in isolation. In the latter model, security is built top-down, starting from the highest level of abstraction and introducing only the essential details. This abstraction when instantiated with an appropriate distinguishability notion allows to prove composable security, that is, that the protocol is guaranteed to remain secure even when integrated with other protocols to construct a more complex cryptographic system.

2.2.1 Stand-alone security model

As mentioned in the introduction to the section, the quantum stand-alone security model is a semantic security framework where the definitions are built from the lowest levels of abstraction.

In this section we will give a short introduction to the model, starting from the parties participating in the protocol, including the ideal functionality and the environment, which are modelled as interactive machines. We will then define precisely what the ideal and real worlds for a run of a protocol are, and when these are indistinguishable, making the protocol secure.

Quantum Interactive Machines (QIM). A quantum interactive machine $\mathbf{A} = \{A_\lambda\}_{\lambda \in \mathbb{N}}$ is a sequence of quantum circuits A_λ, indexed by the security parameter λ, working on an input, output and network register. Two machines can interact by sharing their network register while they are activated alternately. A (two-party) protocol $\Pi = (\mathbf{A}, \mathbf{B})$ is a pair of QIM. We denote by $\mathbf{A} \leftrightsquigarrow \mathbf{B}$ the sequence of quantum maps (indexed by $\lambda \in \mathbb{N}$) representing the interaction between A_λ and B_λ: Namely this map takes as input a quantum state on two registers S_A and S_B, provides to A_λ (resp. B_λ) the

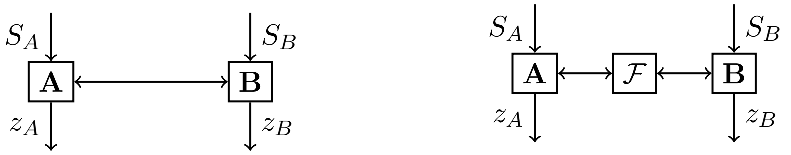

(a) Representation of $\mathbf{A} \leftrightsquigarrow \mathbf{B}$.

(b) Representation of $\mathbf{A} \overset{\mathcal{F}}{\leftrightsquigarrow} \mathbf{B}$.

Figure 2.1: Interaction of two quantum interactive machines **A** and **B**, with and without functionality $\mathcal{F}$.

input S_A (resp. S_B), lets A_λ and B_λ interact and outputs at the end of the interaction the two registers containing the outputs of A_λ and B_λ. We might also write $z \leftarrow \mathsf{OUT}_{\mathbf{B}}(\mathbf{A}(x) \leftrightsquigarrow \mathbf{B}(y))$ instead of $(\_, z) \leftarrow (\mathbf{A}(x) \leftrightsquigarrow \mathbf{B}(y))$ to denote the output of the party **B**. A protocol is said to be *poly-time* if all the parties run in polynomial time. A functionality $\mathcal{F}$ (having no input) is a QIM playing the of a trusted third party interacting with all parties: for two QIM **A** and **B**, we similarly denote with $\mathbf{A} \overset{\mathcal{F}}{\leftrightsquigarrow} \mathbf{B}$ the quantum map that forwards the two input registers to **A** and **B** and that returns their outputs after letting both of them interact (only) with $\mathcal{F}$, as pictured in figure 2.1. We might also provide access to oracles H (QIM that answer queries to functions, e.g. a random oracle), in which case we will either denote it as $\mathbf{A} \overset{H}{\leftrightsquigarrow} \mathbf{B}$ (in this case H is the functionality that answers queries and forwards other messages). Moreover, for two sequences of quantum circuits $\mathbf{A} = \{A_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathbf{B} = \{B_\lambda\}_{\lambda \in \mathbb{N}}$, we also naturally define their sequential composition as $\mathbf{AB} := \{A_\lambda B_\lambda\}_{\lambda \in \mathbb{N}}$.

Adversaries. An adversary $\mathcal{A}$ is also a QIM, able to corrupt parties. In the quantum stand-alone model we only consider *static* adversaries, meaning that the set of corrupted parties is chosen at the beginning of the protocol. We denote by $\hat{\mathbf{A}}$ the adversary that corrupts **A** (similarly $\hat{\mathbf{B}}$ would corrupt **B**), thus in a two-party protocol $\mathcal{A} \in \{\hat{\mathbf{A}}, \hat{\mathbf{B}}\}$. We define $\Pi \leftrightsquigarrow \mathcal{A}$ as the quantum map obtained when the protocol Π is run in the presence of the adversary $\mathcal{A}$: Notably, when $\Pi = (\mathbf{A}, \mathbf{B})$, we have $(\Pi \leftrightsquigarrow \hat{\mathbf{A}}) = (\hat{\mathbf{A}} \leftrightsquigarrow \mathbf{B})$ and $(\Pi \leftrightsquigarrow \hat{\mathbf{B}}) = (\mathbf{A} \leftrightsquigarrow \hat{\mathbf{B}})$.

Distinguisher. A distinguisher **Z**, a QIM outputting a single classical bit, is in charge of distinguishing a real realization of the protocol from an ideal one.

Definition 2.2.1 (Indistinguishable quantum maps)**.** *Two sequences of quantum maps* $\mathbf{X} = \{X_\lambda\}_{\lambda\in\mathbb{N}}$ *and* $\mathbf{Y} = \{Y_\lambda\}_{\lambda\in\mathbb{N}}$ *are said to be* computationally *(resp.* statistically*)* indistinguishable*, denoted* $\mathbf{X} \approx_c \mathbf{Y}$ *(resp.* $\mathbf{X} \approx_s \mathbf{Y}$*), if for any poly-time (resp. unbounded) QIM* $\mathbf{Z} = \{Z_\lambda\}_{\lambda\in\mathbb{N}}$ *and any sequence of advice* $\sigma = \{\sigma_\lambda\}_\lambda$*,*

$$\Big|\Pr(Z_\lambda(X_\lambda \otimes I)\sigma_\lambda = 1) - \Pr(Z_\lambda(Y_\lambda \otimes I)\sigma_\lambda = 1)\Big| \leq \mathsf{negl}(\lambda)\,.$$

Note that the initial advice σ is an input to the two quantum maps and possibly provides side information to the distinguisher. We refer to the joint pair of distinguisher and advice $(\mathbf{Z}, \sigma)$ as the *environment.*

Real and ideal worlds. In order to *simulate* the real world with an ideal world, we replace any honest party $\mathbf{A}$ by an idealized party $\tilde{\mathbf{A}}$ that honestly interacts with $\mathcal{F}$, typically trivially interacting with $\mathcal{F}$ by forwarding the inputs and outputs to/from $\mathcal{F}$, and we write $\tilde{\Pi} := (\tilde{A}, \tilde{B})$ to denote this dummy protocol. To simulate the corrupted parties in the ideal world, we introduce a special kind of adversary $\mathsf{Sim}_\mathcal{A}$ called a *simulator*, that must corrupt the same party as the adversary $\mathcal{A}$.

Definition 2.2.2. *Let* $\Pi = (\mathbf{A}, \mathbf{B})$ *be a two-party protocol and let* $\mathcal{A}$ *be a static adversary. Let* $\mathsf{Sim}_\mathcal{A}$ *be a simulator,* $\sigma = \{\sigma_\lambda\}_{\lambda\in\mathbb{N}}$ *a sequence of quantum states and* $\mathbf{Z}$ *a distinguisher. We denote by* $\mathsf{REAL}^{\sigma}_{\Pi,\mathcal{A},\mathbf{Z}} := \mathbf{Z}((\Pi \leftrightsquigarrow \mathcal{A}) \otimes I)\sigma$ *the sequence of binary random variables output by the distinguisher* $\mathbf{Z}$ *at the end of a real run of the protocol* Π *in presence of an adversary* $\mathcal{A}$*. We denote similarly* $\mathsf{IDEAL}^{\sigma,\mathcal{F}}_{\tilde{\Pi},\mathsf{Sim}_\mathcal{A},\mathbf{Z}} := \mathbf{Z}((\tilde{\Pi} \overset{\mathcal{F}}{\leftrightsquigarrow} \mathsf{Sim}_\mathcal{A}) \otimes I)\sigma$ *the sequence of binary random variables output by the distinguisher* $\mathbf{Z}$ *at the end of an ideal run of the protocol* $\widetilde{\Pi}$ *where the simulator can corrupt the same dummy parties as the adversary* $\mathcal{A}$ *interacting with the ideal functionality* $\mathcal{F}$*.*

The real and ideal executions of the protocol $\Pi = (\mathbf{A}, \mathbf{B})$, when a malicious Berain $\mathcal{A} = \mathbf{B}$ is considered, are depicted in figure 2.2.

Definition 2.2.3. *Let* $\mathcal{F}$ *be a poly-time two-party functionality and* Π *be a poly-time two-party protocol. We say that* Π computationally quantum-stand-alone*, denoted* $\mathsf{C-QSA}$*, (resp.* statistically *quantum-stand-alone, denoted* $\mathsf{S-QSA}$*) realizes* $\mathcal{F}$ *if for any poly-time (resp. unbounded) adversary* $\mathcal{A}$ *there is a poly-time (in the time taken by* $\mathcal{A}$*) simulator* $\mathsf{Sim}_\mathcal{A}$ *such that for any poly-time (resp. unbounded) environment* $(\mathbf{Z}, \sigma)$*,* $\mathsf{REAL}^{\sigma}_{\Pi,\mathcal{A},\mathbf{Z}} \approx \mathsf{IDEAL}^{\sigma,\mathcal{F}}_{\tilde{\Pi},\mathsf{Sim}_\mathcal{A},\mathbf{Z}}$.

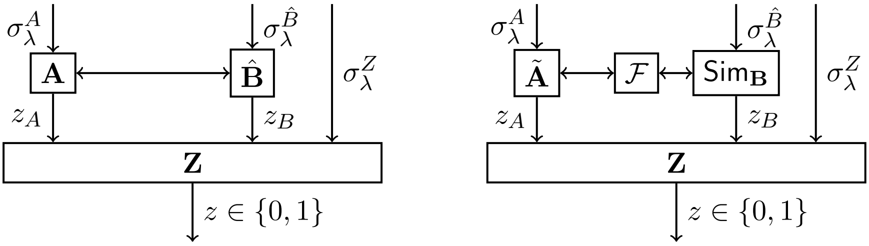

(a) Real-world.

(b) Ideal-world.

Figure 2.2: Real world and ideal world executions when Berain is malicious.

Note that in a multi-party protocol, the adversary $\mathcal{A}$ can corrupt different sets of parties, in such case we will explicitly state this set.

Hybrid models. In order to prove a realization of a functionality $\mathcal{F}$ securely, it is sometimes necessary to make extra assumptions, for example the existence of another (unspecified) protocol realizing a more primitive functionality $\mathcal{G}$. We denote by $\Pi^\mathcal{G}$ a protocol where each party can interact with a trusted party running $\mathcal{G}$, and we say that Π realizes the functionality $\mathcal{F}$ under the assumption $\mathcal{G}$ or that we are in the $\mathcal{G}$-*hybrid* model. If we do not specify how the primitive $\mathcal{G}$ is realized but instead assume its realization is possible, we call this a *setup assumption*. Importantly, if a protocol realizes $\mathcal{G}$ and if a $\mathcal{G}$-hybrid protocol realizes $\mathcal{F}$, then combining both protocols in the natural way gives a protocol realizing $\mathcal{F}$.

2.2.2 Abstract cryptography

Instead of building cryptography from the bottom-up as in the stand-alone security model, we can also start from the highest level of abstraction and define the essential properties we require from our cryptographic tasks.

We follow here the presentation by Portmann [Por17], adapting it to draw parallels with the stand-alone model. Instead of looking at the participants, the abstract cryptography (AC) framework focuses in interaction through the ideal functionality, thus we start by modelling these interactions. We then define what the ideal and real worlds for a run of a protocol are, and when these are indistinguishable, making the protocol secure. The main difference in practice is that in the AC framework the distinguisher is allowed to perform any operation allowed by quantum mechanics and in particular,

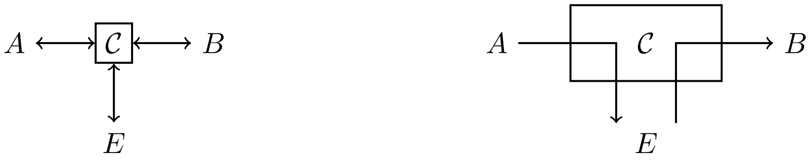

(a) General resources $\mathcal{C}$ accessible to 3 parties.

(b) Insecure channel $\mathcal{C}$ accessible to 3 parties, here E is the adversary.

Figure 2.3: Resources.

the environment is not reduced to interacting with the parties in a sequential manner.

Interactions between parties. In an n-player setting, a *resource* is an object with n interfaces; it allows the players to input and receive messages. We will denote resources by squares, and inputs/outputs from the interfaces by lines intersecting with the squares, see figure 2.3. If two resources $\mathcal{C}$ and $\mathcal{K}$ are available to the players, we write $\mathcal{C}||\mathcal{K}$ for the parallel composition of the resources: the resources are simultaneously accessible to the players in any arbitrary order, thus in particular, the order of composition is irrelevant and $\mathcal{C}||\mathcal{K} = \mathcal{K}||\mathcal{C}$. Some of these resources will indeed be the ideal functionalities discussed in Section 2.2.1.

Honest parties. A *converter* models the local operations that the players can perform through their interfaces. We will denote converters by squares with rounded corners. If a converter σ is connected to the interface i of the resource $\mathcal{C}$, we write $\sigma_i\mathcal{C}$ (or equivalently $\mathcal{C}\sigma_i$), see figure 2.4a. A *protocol* is defined by a set of converters: One for each honest player, and are denoted by π.

Distinguisher. On the one hand, an adversary is allowed to perform any operation allowed by quantum mechanics. On the other hand, to prove security guarantees it is not enough to show indistinguishability in the presence of adversaries, we also need to emulate the presence of no adversary. We do this with a special type of converter, called *filters*, that emulate an honest behaviour, for example without any tampering but still with noise. We call *filtered resources* the pair of resource $\mathcal{C}$ and filter $\Diamond_E$, denoted $\mathcal{C}_\Diamond = (\mathcal{C}, \Diamond_E)$. See figure 2.4b for a noisy filter $\#_E$ replacing an adversary.

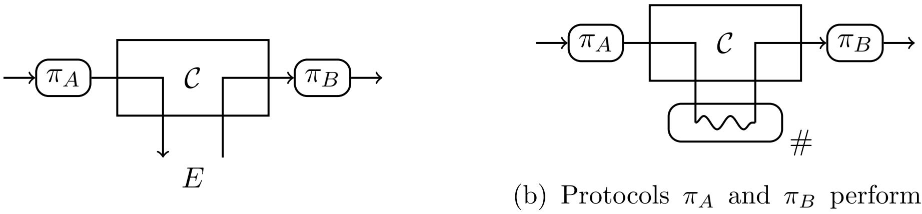

(a) Protocols π_A and π_B performed by two honest players and adversary E, with access to insecure channel $\mathcal{C}$.

(b) Protocols π_A and π_B performed by two honest players with access to insecure channel $\mathcal{C}$, the adversary is replaced by a noisy filter $\#$.

Figure 2.4: Converters and filters.

Real and ideal worlds. In an ideal run of the protocol the honest parties interact with the ideal resource and the adversary is replaced by an honest filter, which might still induce noise. For a real run of the protocol the honest parties run their protocols and interact through the untrusted resources, with which the adversary also interacts performing any operation allowed by quantum mechanics.

We say that a protocol is secure when the interactions of the adversary with the real resource can be simulated by interactions with an ideal resource. In the AC framework, we call this *constructing* an ideal resource $\mathcal{S}$ from a real resource $\mathcal{C}$ via the protocol π. We include the formal definition below.

Definition 2.2.4 (Cryptographic security)**.** *We say that the protocol $\pi_{AB} = (\pi_A, \pi_B)$ constructs the filtered resource $\mathcal{S}_\Diamond$ from $\mathcal{C}_\#$ within (ε, δ), denoted $\mathcal{C}_\# \xrightarrow{\pi_{AB},(\varepsilon,\delta)} \mathcal{S}_\Diamond$, if the following two conditions hold:*

i. ***Correctness:*** *In presence of no malicious player, the filtered resources are ε-close to each other*

$$d(\pi_{AB}\mathcal{C}_\#, \mathcal{S}_\Diamond) \leq \varepsilon.$$

ii. ***Security:*** *In the presence of an adversary, there exists a simulator σ_E, δ-close to the real protocol*

$$d(\pi_{AB}\mathcal{C}, \sigma_E\mathcal{S}) \leq \delta.$$

Here the distance d is the supremum over the set of all possible distinguishers allowed by quantum mechanics, which reduces to the diamond norm if the distinguisher only has access to one input port. If the filtered resources $\mathcal{S}_\Diamond$ and $\mathcal{C}_\#$ are clear from the context, we say that π_{AB} is (ε, δ)-secure, or ε-correct and δ-secure.

We differ from the original definition of cryptographic security in [MR11], where security is defined as the maximum of the two values ε and δ, because these parameters have independent meanings that are interesting to study separately. The ε in item i. refers to the *correctness* of the protocol. That is, the probability that the protocol running on a noisy channel without adversary will be distinguishable from an ideal channel. The δ in item ii. is the usual *security* in presence of an adversary.

The main interest in separating correctness and security in the above definition is that the serial composition of two protocols respects the additivity of both parameters separately. We prove this property in the following theorem.

Theorem 2.2.5 (Serial composition security)**.** *Let the protocols π and π' construct $\mathcal{S}_\Diamond$ from $\mathcal{R}_\#$ and $\mathcal{T}_\Box$ from $\mathcal{S}_\Diamond$ within (ε, δ) and (ε', δ') respectively, i.e.*

$$\mathcal{R}_\# \xrightarrow{\pi,(\varepsilon,\delta)} \mathcal{S}_\Diamond \quad \textit{and} \quad \mathcal{S}_\Diamond \xrightarrow{\pi',(\varepsilon',\delta')} \mathcal{T}_\Box \,. \tag{2.11}$$

Then the serial composition $\pi'\pi$ constructs $\mathcal{T}_\Box$ from $\mathcal{R}_\#$ within $(\varepsilon+\varepsilon', \delta+\delta')$,

$$\mathcal{R}_\# \xrightarrow{\pi'\pi,(\varepsilon+\varepsilon',\delta+\delta')} \mathcal{T}_\Box \,. \tag{2.12}$$

Proof. The statement follows directly from the triangle inequality. For ε-correctness we have that

$$d(\pi'\pi\mathcal{R}_\#, \mathcal{T}_\Box) \leq d(\pi'\pi\mathcal{R}_\#, \pi'\mathcal{S}_\Diamond) + d(\pi'\mathcal{S}_\Diamond, \mathcal{T}_\Box) \leq d(\pi\mathcal{R}_\#, \mathcal{S}_\Diamond) + \varepsilon' \leq \varepsilon + \varepsilon' \,. \tag{2.13}$$

Similarly for δ-security, the composed converter $\sigma'\sigma$ is a converter for the composition since

$$\begin{align} d(\pi'\pi\mathcal{R}, \sigma'\sigma\mathcal{T}) &\leq d(\pi'\pi\mathcal{R}, \pi'\sigma\mathcal{S}) + d(\pi'\sigma\mathcal{S}, \sigma'\sigma\mathcal{T}) \tag{2.14}\\ &\leq d(\pi\mathcal{R}, \sigma\mathcal{S}) + d(\pi'\mathcal{S}, \sigma'\mathcal{T}) \tag{2.15}\\ &\leq \delta + \delta' \,, \tag{2.16} \end{align}$$

where we used commutativity of converters $\alpha\beta\mathcal{C} = \beta\alpha\mathcal{C}$ and the pseudo-metric property $d(\alpha\mathcal{C}, \alpha\mathcal{C}') \leq d(\mathcal{C}, \mathcal{C}')$, see [Mau11]. □

Chapter 3

Port-Based State Preparation

This chapter is based on the article *Port-Based State Preparation and Applications* [MS24], and is joint work with Florian Speelman.

A conclusion section has been added and the remaining sections have been edited for style and typographical consistency.

In this chapter we introduce a new variant of quantum teleportation where the sender holds a classical description of the target state, which we call *Port-Based State Preparation*, and in Section 3.3 we analyse the resource requirements for its implementation. In Section 3.4 we give a possible application for designing universal quantum processors with classical inputs, and study its implementation costs. In Section 3.5 we prove the maximum achievable optimality of both tasks. We conclude that these known-state variants of teleportation and computation respectively, require exponentially less pre-shared entanglement, in the error dependence, per qubit than their unknown-state analogues.

3.1 Introduction

Quantum teleportation [BBC$^+$93] is a fundamental task in quantum information processing: Two parties, Anboto and Berain, can implement a quantum channel transmitting a qubit by means of a classical channel and one entangled quantum bit shared between the two parties. Although an infinitely-precise description of a quantum state would require an infinite number of classical bits to describe, quantum teleportation allows an observer who cannot access the infinite information in a quantum state to send this information only with 2 bits of classical communication and an entangled qubit.

Standard teleportation is a protocol where the transmitted quantum state is unknown, but what if Anboto knows which quantum state she would like to transmit? Of course, she could first create the state and then use a standard teleportation protocol, but this extra knowledge might open up the option for other protocols. This variant of teleportation is called remote state preparation (RSP) [BDS$^+$01], and can indeed be performed more efficiently; for instance, asymptotically only one bit of classical communication is necessary to transmit a qubit.

Port-based teleportation [IH08] (PBT) is a variant of teleportation where the receiver's correction operation is very simple; after receiving classical information from Anboto, Berain simply picks one of the subsystems, or *ports*, and traces out the rest of his part of the entangled resource state. This form of the correction operation allows Berain to perform quantum operations on the output of the protocol before receiving the correct port — this property makes PBT a useful primitive for various tasks in quantum information, including proving bounds on channel discrimination [PLLP19], universal programmable quantum processors [NC97, YRC20], constructing protocols for instantaneous non-local quantum computation [BK11], transposing and inverting unitary operations [QDS$^+$19a, QDS$^+$19b], storage and retrieval of unitary channels [SBZ19] and relating communication complexity with Bell nonlocality [BCG$^+$16].

The PBT primitive is powerful but not very efficient, requiring a number of ports scaling polynomially with the dimension of the transmitted state (i.e., exponential in the number of qubits transmitted) and the tolerated error. Given that knowing the teleported quantum state makes RSP more efficient than standard teleportation, we will study the following question:

What is the complexity of port-based teleportation of a known state?
And what would be the uses of such primitive?

Our contributions. In this chapter, we introduce and study a variant of quantum teleportation, port-based state preparation (PBSP); where the target qubit is known to Anboto, but is still unknown to Berain, and Berain cannot perform any correction operation other than discarding ports. We can consider probabilistic and deterministic variants of PBSP, where the former either constructs the desired state perfectly or signals a failure, and the latter always constructs a (necessarily approximate) state. A protocol for this task is present in the RSP literature, denoted as the 'column method' [BHL+05], but as far as we are aware we are the first ones describing the general task explicitly and studying the relations between the dimension of the entangled resource state, the number of ports and the error of the outcome state.

Since Anboto knows the state she intends to send, her measurement can depend on it, which turns out to be enough for the error to scale inverse-exponentially in the number of ports, both for the deterministic and the probabilistic PBSP. Moreover, given a d dimensional input state, when N independent maximally entangled states are taken as resource states, we design measurements in Section 3.3 that achieve the same success probability p and worst-case fidelity F_{wc},

$$p = F_{\mathrm{wc}} = 1 - \left(1 - \frac{1}{d}\right)^N, \tag{3.1}$$

for the probabilistic and deterministic PBSP respectively. Moreover, in Section 3.5 we show that this is the optimal achievable error for both the probabilistic and deterministic PBSP when maximally entangled states are used as resource.

Motivated by the multiple applications of PBT, we also introduce and study a known-state variant of a universal programmable processor [NC97], which we call universal programmable hybrid processor (UPHP); a quantum machine that given a description of a quantum state, outputs a state approximately close to the state with a chosen unitary applied to it. The name is meant to evoke a combination between classical and quantum universal processors; classical computers can compute any function on input data by programming a universal gate array, and similarly we can program quantum gate arrays to perform arbitrary unitary operations on quantum data.

Since there are no correction operations, the PBSP protocol allows the receiver to apply any quantum operation in their registers before receiving the classical message from the sender. We show in Theorem 3.4.2 that the memory dimension m needed for the UPHP processor build from PBSP scales

	Existence	
PBT	$f \geq 1 - \frac{d(d-1)}{N}$	[IH08]
PBSP	$F_{\text{wc}} = 1 - \left(1 - \frac{1}{d}\right)^N$	Thm. 3.3.2

	Optimality	
PBT	$f^* \leq 1 - \frac{d(d-1)}{8N^2} + O\left(\frac{1}{N^3}\right)$	[CLM⁺21]
	$f^* \leq K^{\frac{\log(d)}{d}}\left(2N + \frac{2}{3}\right)$	[KPPG19]
PBSP	$1 - h(\sqrt{1 - F^*_{\text{wc}}}) \leq \frac{\log(d)}{d} 4N$	Thm. 3.5.2
	$F^*_{\text{wc,EPR}} \leq 1 - \left(1 - \frac{1}{d}\right)^N$	Thm. 3.5.4

Table 3.1: Average fidelity f (see Definition 3.2.3) and worst-case fidelity (see Definition 3.2.2) of teleporting a qudit with N ports for deterministic PBT and PBSP.

in terms of relative error ε as:

$$m \leq \left(\frac{1}{\varepsilon}\right)^{4d\ln(d)}. \tag{3.2}$$

Once again, this shows an exponential separation between the classical-quantum hybrid and fully-quantum universal processors. Finally, in Theorem 3.5.1 we bound the optimal achievable memory dimension for UPHPs by

$$m \geq 2^{\frac{(1-h(\varepsilon))}{2}d}, \tag{3.3}$$

which follows from the impossibility of quantum random access codes (QRACs) coming from Nayak's bound, see Theorem 3.2.5.

Relation to the literature. Our results show a fundamental difference between having an unknown versus known state on Anboto's side. The PBSP protocol achieves an exponential scaling of the error in terms of the number of ports, in contrast to the PBT results, where linear upper bounds exist for optimal scaling for any resource state, see tables 3.1 and 3.2. For deterministic PBT it is known that the pretty good measurement (PGM) achieves the optimal fidelity both for maximally entangled resource states and when the resource states are optimized [Led22], but this is not the case for probabilistic PBT [PG13, MSSH18]. Although probabilistic and deterministic PBT

	Existence	
PBT	$p = \frac{N}{N-1+d^2}$	[SSMH17]
PBSP	$p = 1 - \left(1 - \frac{1}{d}\right)^N$	Thm. 3.3.1

	Optimality	
PBT	$p^* \leq \frac{N}{N-1+d^2}$	[PG13, MSSH18]
PBSP	$p^*_{\text{EPR}} \leq 1 - \left(1 - \frac{1}{d}\right)^N$	Thm. 3.5.3

Table 3.2: Probability p of teleporting a qudit with N ports for probabilistic PBT and PBSP.

behave differently not only in terms of achievable error but also the measurements involved, for both PBSP tasks we obtain exponentially decreasing error with maximally mixed resource states and practically the same measurement. We include a summary of the comparison between PBT and PBSP results in tables 3.1 and 3.2, note that the optimality for PBSP is only discussed for maximally entangled resources, which we denote by an EPR subscript.

Along with the definition of UPQPs, Nielsen and Chuang proved that they are impossible to perform perfectly [NC97], and many have afterwards worked on closing the gap between existence and impossibility results from different directions; e.g. studying the possible overlap between two program states [HZB06], through epsilon-nets [Maj18], or with type constants of Banach spaces [KPPG19]. Finally, a recent result by Yang, Renner and Chiribella [YRC20] closed the gap with a tight bound on the memory dimension, which scales d^2-inverse-exponentially with the error, where d is the dimension of the unitary operation. In this regard, the existence result of UPHP we provide scales $d\log(d)$-inverse-exponentially with the error, implying that UPHPs are an exponentially easier task to perform than UPQPs. Moreover, our lower bound obtained using the well-known Nayak's bound for QRACs also works for UPQPs and although weaker than the optimal one, it improves on the previous best known one by Kubicki, Palazuelos and Perez-García [KPPG19]. We include a summary of the comparison between UPQPs and UPHPs results in table 3.3.

	Upper bounds		Lower bounds	
UPQP	$\left(\frac{1}{\varepsilon}\right)^{d^2}$	[KPPG19]	$2^{\frac{(1-\varepsilon)}{12}d-\frac{2}{3}\log(d)}$ $\frac{(1-\varepsilon)}{4}d$ if unitary	[KPPG19]
	$\left(\frac{1}{\varepsilon}\right)^{\frac{d^2-1}{2}}$	[YRC20]	$\left(\frac{1}{\varepsilon}\right)^{\frac{d^2-1}{2}-\alpha}$, $\forall\alpha > 0$	[YRC20]
UPHP	$\left(\frac{1}{\varepsilon}\right)^{4d\ln(d)}$	Thm. 3.4.2	$2^{\frac{d}{2}(1-h(2\varepsilon))}$	Thm. 3.5.1

Table 3.3: Memory dimension m necessary for universally programming any unitary in d dimensions, the set $\mathrm{U}(d)$, with UPQPs and UPHPs up to error ε.

3.2 Preliminaries

Although the general quantum-information notation can be found in Chapter 2, here we will introduce some concepts and prior results that we need only in this chapter. In particular, we describe the standard port-based teleportation and universal programmable quantum processors tasks.

3.2.1 Notation

We define *maximally entangled* qudits as

$$\left|\phi_d^+\right\rangle := \frac{1}{\sqrt{d}}\sum_{x=1}^{d}|x\rangle \otimes |x\rangle \,. \tag{3.4}$$

Given $A := A_1 \dots A_n$, and $x \in [N]$, we denote by $A_{\bar{x}}$ all the elements in A other than the x-th, this is, $A_{\bar{x}} := A_1, \dots, A_{x-1}, A_{x+1}, \dots, A_N$.

We use ψ to denote the classical description of a quantum state $|\psi\rangle$. Because our results are of information-theoretic nature, we will assume the classical description to be of arbitrarily-high precision representation of $|\psi\rangle$ unless specified otherwise.

We denote by $F(\rho, \sigma)$ the fidelity between two states

$$F(\rho,\sigma) = \left\|\sqrt{\rho}\sqrt{\sigma}\right\|_1^2 = \mathrm{Tr}\left[\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}\right]^2 . \tag{3.5}$$

If one of the states is pure, say $\rho = |\psi\rangle\!\langle\psi|$, then the fidelity is the overlap of the states

$$F(\rho,\sigma) = \mathrm{Tr}\left[\sqrt{|\psi\rangle\!\langle\psi|\sigma|\psi\rangle\!\langle\psi|}\right]^2 = \langle\psi|\sigma|\psi\rangle \,. \tag{3.6}$$

Moreover, by the Fuchs-van de Graaf inequalities we have the following relation between the fidelity and distance between two states,

$$1 - \sqrt{F(\rho,\sigma)} \leq T(\rho,\sigma) \leq \sqrt{1 - F(\rho,\sigma)}\,. \tag{3.7}$$

Definition 3.2.1 (Entanglement fidelity)**.** *Given two channels* $\Phi^1_A, \Phi^2_A \in \mathrm{CPTP}(\mathcal{H}_A)$*, we define their* entanglement fidelity *as*

$$F(\Phi^1_A, \Phi^2_A) := F((\Phi^1_A \otimes I_B)|\phi^+_{AB}\rangle\langle\phi^+_{AB}|, (\Phi^2_A \otimes I_B)|\phi^+_{AB}\rangle\langle\phi^+_{AB}|)\,.$$

Definition 3.2.2 (Worst-case fidelity)**.** *Given* $\Phi^1_A, \Phi^2_A \in \mathrm{CPTP}(\mathcal{H}_A)$*, we define their* worst-case fidelity *as*

$$F_{\mathrm{wc}}(\Phi^1_A, \Phi^2_A) := \inf_{|\psi\rangle} F((\Phi^1_A \otimes I_B)|\psi_{AB}\rangle\langle\psi_{AB}|, (\Phi^2_A \otimes I_B)|\psi_{AB}\rangle\langle\psi_{AB}|)\,.$$

We denote by $F^*(\Phi, \mathcal{I})$ (respectively $F^*_{\mathrm{wc}}(\Phi, \mathcal{I})$) the optimal achievable entanglement (respectively worst-case) fidelity between a channel Φ and the identity channel $\mathcal{I}$, we will call this the entanglement (respectively worst-case) fidelity of the channel Φ and denote it by F^* (respectively F^*_{wc}) if the protocol we are talking about is clear from the context.

Definition 3.2.3 (Average fidelity)**.** *Given a channel* $\Phi \in \mathrm{CPTP}(\mathcal{H}_A)$*, we define its* average fidelity *as*

$$f(\Phi) := \int_\psi \langle\psi|\Phi(|\psi\rangle\langle\psi|)|\psi\rangle \mathrm{d}\psi\,.$$

We denote by $f^*(\Phi)$ the optimal achievable average of the channel Φ, and denote it by f^* if the protocol we are talking about is clear from the context. There is a well-known relation between average fidelity and entanglement fidelity of any channel, see [HHH99, Proposition 1]:

$$f(\Phi) = \frac{F(\Phi, \mathcal{I})d + 1}{d + 1}\,. \tag{3.8}$$

3.2.2 Quantum Random Access Codes

A fundamental impossibility result of quantum information theory is Holevo's theorem, i.e. no more than n bits of *accessible* information can be transmitted by transferring n qubits. Intuitively, Holevo's bound comes from the necessity of choosing a measurement, which will possibly destroy the information that could have been obtained by another measurement. Quantum random access codes (QRACs) were introduced by Ambainis, Nayak, Ta-Shma, and Vazirani [ANTV99] with the hope of circumventing Holevo's theorem. The idea was to hand some measurements to the receiver —as many measurements as bits— such that each measurement would reveal at least one of the bits with a certain fixed probability of success. Note that the existence of QRACs would not immediately violate Holevo's bound: as the measurements do not commute, they cannot be applied subsequently with the same success probability.

Definition 3.2.4 (QRAC). *A (n, m, p)-quantum random access code is a mapping of n classical bits into m qubits ($f : \{0,1\}^n \to \mathbb{C}^{2^m}$, $x \mapsto \rho_x$), along with a set of measurements $(M_1^0, M_1^1), \ldots, (M_n^0, M_n^1)$ satisfying, for all $x \in \{0,1\}^n$ and $i \in [n]$,*

$$\operatorname{Tr}(M_i^{x_i} f(x)) \geq p\,.$$

Unfortunately, QRACs do not improve much on Holevo's bound, as shown in [Nay99]. In particular, although an n-qubit quantum state is a vector in a 2^n-dimensional complex space, it —perhaps surprisingly— cannot encode an amount of information that is more than linear in n; or at least not in an accessible way.

Theorem 3.2.5 (Nayak's bound)**.** *Any (n, m, p)-QRAC must satisfy*

$$m \geq n(1 - h(p))\,,$$

where h is the binary entropy function.

3.2.3 Port-Based Teleportation

On the other side of the story, we can ask if the converse of such a statement is true, that is, what the number of bits we need to specify a qubit is. Intuitively, since we can only approximate a complex number with a finite number of bits, this should be impossible, but teleportation tasks circumvent this limitation by making use of entanglement as a resource.

Here we are interested in a particular teleportation task called port-based teleportation (PBT), introduced by Ishizaka and Hiroshima [IH08], where Berain's computational power is restricted to tracing out registers. Note that in this situation the only meaningful classical communication is communicating a port, thus the classical communication cost is completely determined by the number of registers that Berain holds, in other words, the dimension of Berain's part of the resource state.

Protocol 1 (d, N, p) Probabilistic Port-Based Teleportation

Setting. Both Anboto and Berain each have N qudits in $A_1, A_2, \dots, A_N$ and $B_1, B_2, \dots, B_N$ respectively, which are denoted by A and B as a whole. They share an arbitrary entangled resource in AB, and Anboto holds an additional input register A_0. There is a single round of classical communication from Anboto to Berain. Berain's correction operations are restricted to tracing out registers.

Inputs. A d-dimensional quantum state in Anboto's register A_0.

Goal. Anboto's input state to be teleported to an arbitrary output port of Berain, with probability of success p.

The protocol:

```
1: Anboto performs a $N + 1$ outcome POVM measurement on her $d$-dimensional input state and her half of the $N$ shared resource states.
2: Anboto communicates either a port $x \in [N]$ or abort $x = 0$ to Berain with $\log(N + 1)$ bits of classical communication.
3: If $x = 0$ Berain aborts, else he traces out all but register $B_x$.
```

We can distinguish two types of PBT tasks; probabilistic and deterministic ones. In the probabilistic case, Anboto performs a $N+1$ outcome POVM $\mathcal{M} = \{M_x\}_{x=0}^{N}$ such that only outcomes $x = 1, \dots, N$ refer to a successful transmission of the state, in this case the state in port x has perfect fidelity with respect to the state intended, and outcome $x = 0$ refers to a failure in the transmission, therefore Berain aborts. In this setting it is natural to study the performance of the protocol by studying the success probability. To avoid confusion, we will denote these as (d, N, p)-PBT protocols, where p denotes the probability of success of the protocol.

Deterministic PBT does not consider the option of failure, instead Anboto performs a N outcome POVM, and always transmits a port. It is clear that such a protocol must be approximate, meaning that the outcome state cannot always have perfect fidelity with the intended one, hence we study the performance of the protocol in terms of the entanglement fidelity of the channel.

For protocols with a covariance property, as is the case of PBT, both the entanglement fidelity F and the diamond norm ε are equivalent [YRC20, Appendix B, Lemma 2],

$$F = 1 - \frac{\varepsilon}{2}, \tag{3.9}$$

Protocol 2 (d, N, F) Deterministic Port-Based Teleportation

Setting. Both Anboto and Berain each have N qudits in $A_1, A_2, \ldots, A_N$ and $B_1, B_2, \ldots, B_N$ respectively, which are denoted by A and B as a whole. They share an arbitrary entangled resource in AB, and Anboto holds an additional input register A_0. There is a single round of classical communication from Anboto to Berain. Berain's correction operations are restricted to tracing out registers.

Inputs. A d-dimensional quantum state in Anboto's register A_0.

Goal. Anboto's input state to be teleported to an arbitrary output port of Berain, with fidelity F.

The protocol:

```
1: Anboto performs a $N$ outcome POVM measurement on her $d$-dimensional
   input state and her half of the $N$ shared resource states.
2: Anboto always communicates a port $x \in [N]$ to Berain with $\log(N)$ bits
   of classical communication.
3: Berain traces out all but register $B_x$.
```

In the rest of this chapter we will talk about probability of success and entanglement fidelity as measures of accuracy.

In particular, any probabilistic (d, N, p)-PBT can be used to construct a deterministic one, just by Anboto transmitting a random register index whenever the measurement outcome is failure, i.e. $x = 0$. The achieved entanglement fidelity is

$$F \geq p + \frac{1-p}{d^2}\,. \tag{3.10}$$

There are some known trade-offs between the dimension of the input state d, the number of ports N, and the accuracy of the protocol, depending on the resource states, for both the probabilistic and the deterministic PBT. As all these parameters can be optimized independently, in order to simplify the analysis it is standard in the literature to consider N independent maximally entangled resource states (one per port) and the pretty good measurement, in this case we call the protocol *standard.* We include the known results between resource trade-offs (as far as we are aware) in table 3.4.

	Existence	
Det.	$F^{\text{sta}} \geq 1 - \frac{d^2-1}{N}$	[IH08]
	$\varepsilon^{\text{sta}} \leq \frac{4d^2}{\sqrt{N}}$	[BK11]
	$F \geq 1 - \frac{d^2(d^2-2)-1}{d^2(d^2+N-1)}$	eq. (3.10)
Prob.	$p = \frac{N}{N-1+d^2}$	[SSMH17]

	Optimality	
Det.	$F^* \leq 1 - \frac{1}{4(d-1)N^2} + O\left(\frac{1}{N^3}\right)$	[Ish15]
	$F^* \leq 1 - \frac{d^2-1}{8N^2} + O\left(\frac{1}{N^3}\right)$	[CLM⁺21]
	$F^* \leq K\frac{\log(d)}{d}\left(2N + \frac{2}{3}\right)$	[KPPG19]
Prob.	$p^* \leq \frac{N}{N-1+d^2}$	[PG13, MSSH18]

Table 3.4: Probability p and entanglement fidelity F (see Definition 3.2.1) of teleporting a qudit with N ports for PBT.

3.2.4 Universal Programmable Quantum Processors

An interesting application of PBT is of universal programmable quantum processors (UPQPs). We ideally want to build quantum computers as we do with the classical ones, by building a universal processor independent of the data and programs that we insert, only that now both these elements can be quantum states or unitary. Informally the task is the following: given a unitary, we are asked to apply it in an unknown state, but only after losing access to the unitaries. Therefore, we want to find a way to store the unitary in a quantum memory.

Definition 3.2.6 (UPQP)**.** *A quantum channel* $\mathsf{P} \in \mathrm{CPTP}(\mathcal{H}_D \otimes \mathcal{H}_M)$ *is a d-dimensional Universal Programmable Quantum Processor if for every unitary* $U \in \mathrm{U}(\mathcal{H}_D)$*, there exists a unit vector* $|\phi_U\rangle \in \mathrm{D}(\mathcal{H}_M)$ *such that*

$$\mathrm{Tr}_M[\mathsf{P}(\rho_D \otimes |\phi_U\rangle\langle\phi_U|_M)] = U\rho_D U^\dagger, \quad \text{for all } \rho_D \in \mathrm{D}(\mathcal{H}_D)\,.$$

Along their introduction, Nielsen and Chuang [NC97] proved the impossibility of perfectly implementing UPQPs. The basic idea is that for every unitary we desire to apply, there needs to be an orthogonal memory state, and since for any given dimension there are infinitely many possible unitaries, we need an infinite dimensional memory space to store them. However, the

Upper bounds		Lower bounds	
Error dependency			
$\left(\frac{1}{\varepsilon}\right)^{d^2}$	[KPPG19]	$\frac{d^2}{\varepsilon}$	[Maj18]
$\left(\frac{1}{\varepsilon}\right)^{\frac{d^2-1}{2}}$	[YRC20]	$\left(\frac{1}{\varepsilon}\right)^{\frac{d^2-1}{2}-\alpha}$, $\forall\alpha > 0$	[YRC20]
Dimension dependency in log power			
$d^2 \log(1/\varepsilon)$	[KPPG19]	$\frac{(1-\varepsilon)}{12}d - \frac{2}{3}\log(d)$ $\frac{(1-\varepsilon)}{4}d$ if unitary	[KPPG19]
$d^2 \frac{\log(1/\varepsilon)}{2}$	[YRC20]	$d^2 \frac{\log(1/\varepsilon)}{2} - \beta$, $\forall\beta > \frac{\log(1/\varepsilon)}{2}$	[YRC20]

Table 3.5: Performance of d-dimensional ε-UPQP in terms of memory dimension m.

proof does not hold if we relax the previous definition; that is, for approximate UPQPs, where we only want to implement the desired unitary with approximate accuracy (or with certain probability of success).

Definition 3.2.7 (Approximate UPQP). *A channel* $\mathsf{P} \in \mathrm{CPTP}(\mathcal{H}_D \otimes \mathcal{H}_M)$ *is a d-dimensional ε-Universal Programmable Quantum Processor, if for every unitary* $U \in \mathrm{U}(\mathcal{H}_D)$ *there exists a unit vector* $|\phi_U\rangle \in \mathrm{D}(\mathcal{H}_M)$ *such that*

$$\frac{1}{2}\left\| \mathrm{Tr}_M[\mathsf{P}(\cdot \otimes |\phi_U\rangle\langle\phi_U|_M)] - U(\cdot)U^\dagger \right\|_\diamond \leq \varepsilon\,.$$

However, we cannot do much better for approximate UPQPs. Yang, Renner and Chiribella [YRC20] showed tight upper and lower bounds for the memory dimension that both blows up inverse polynomially with small error ε, and exponentially when increasing the dimension of the input state d, closing a long-lasting open question. We include the known results between resource trade-offs (as far as we are aware) in table 3.5.

3.3 Port-Based State Preparation

Here we introduce a classical-quantum hybrid communication task where Anboto holds a complete classical description of a quantum state, and her goal is for Berain to end up with the quantum state whose description she holds in one of his ports. In order to achieve this, both parties share a finite amount of entangled resources, Anboto performs a measurement on her part of the shared resources and Berain's actions are restricted to tracing out his side of the resources, or 'ports'. We call this port-based state preparation (PBSP).

In a way, the task is a hybrid between remote state preparation (RSP) tasks introduced in the seminal paper by Bennett, DiVicenzo, Shor, Smolin, Terhal and Wootters [BDS$^+$01], and port-based teleportation (PBT). Recall that RSP achieves asymptotic rates of one qubit transmission per bit by allowing Berain to perform any operation on his side, and we are interested in knowing what happens if we restrict Berain to covariant classical operations, as in the case of PBT.

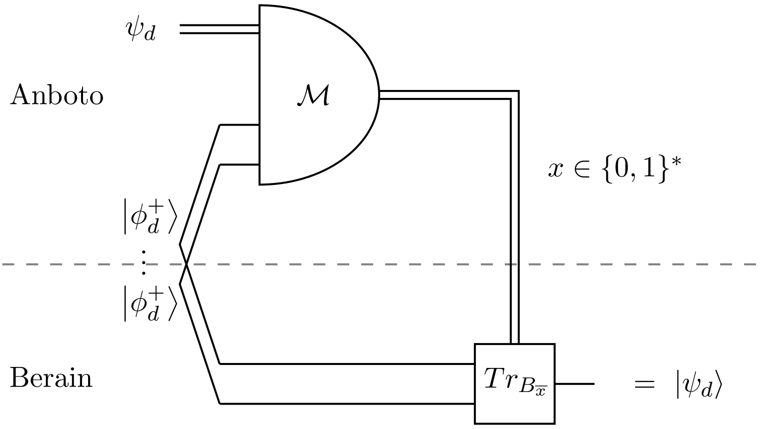

Figure 3.1: Port-Based State Preparation protocol.

We can distinguish between probabilistic and deterministic PBSP tasks. Assuming that both parties share N entangled states, or ports, in probabilistic PBSP (Protocol 3), Anboto's measurement has $N + 1$ possible outcomes; one per port, after which Berain's register will be exactly in the desired state, and one for the failure probability. Naturally we use the protocol's success probability as an accuracy measure.

In deterministic PBSP (Protocol 4), Anboto's measurement has N possible outcomes, one per port, but in this case it is impossible for Berain's register after tracing out to be exactly Anboto's intended state for all the inputs. In this case, we use the fidelity of the outcome state with respect to the input as a measure of success of the protocol.

Protocol 3 (d, N, p) Probabilistic Port-Based State Preparation.

Setting. Both Anboto and Berain each have N qudits in $A_1, A_2, \dots, A_N$ and $B_1, B_2, \dots, B_N$ respectively, which are denoted by A and B as a whole. They share an arbitrary entangled resource in AB. There is a single round of classical communication from Anboto to Berain. Berain's correction operations are restricted to tracing out registers.

Inputs. Anboto receives the classical description of a d-dimensional quantum state.

Goal. Anboto's input state to be prepared in an arbitrary output port of Berain, with probability of success p.

The protocol:

```
1: Anboto performs a $N+1$ outcome POVM measurement on her half of the $N$ shared resource states, which can depend on the classical description of her state.
2: Anboto communicates either a port $x \in [N]$ or abort $x = 0$ to Berain with $\log(N+1)$ bits of classical communication.
3: If $x = 0$ Berain aborts. Else he traces out all but register $B_x$.
```

Remark 3.3.1. *Conventional channel fidelity definitions consider purifications of a given state with respect to a reference space and apply the channel to half of it, as is the case in Definitions 3.2.1 and 3.2.2. Nevertheless the input of our PBSP protocol is a classical description of the state we intend to send, thus it is not clear what it would mean to apply our classical-to-quantum map to half of it. We could instead consider average fidelity as a figure of merit, see Definition 3.2.3, this is*

$$f(\Phi) := \int_\psi \langle\psi|\Phi(\psi)|\psi\rangle \mathrm{d}\psi \,. \tag{3.11}$$

However, this integral is not necessarily well defined. Instead, we define the worst-case fidelity *of a classical-to-quantum map as*

$$F_{\mathrm{wc}}(\Phi) := \inf_{|\psi\rangle\in\mathcal{H}} F(\Phi(\psi), |\psi\rangle) \,. \tag{3.12}$$

Intuitively the covariance-like property of PBSP makes the worst-case fidelity comparable to average fidelity. Clearly, $F_{\mathrm{wc}}(\Phi) \leq f(\Phi)$. Moreover, given a map $\Phi : \{0,1\}^* \to \mathcal{H}_d$, there exists a map $\hat{\Phi} : \{0,1\}^* \to \mathcal{H}_d$ such that $F_{\mathrm{wc}}(\hat{\Phi}) \geq f(\Phi)$. For example, we can define the map $\hat{\Phi}$ by Anboto and Berain *first agreeing to a random d-dimensional unitary $U \sim \sigma_d$ sampled with respect to the Haar-measure, then performing the protocol $\psi \mapsto \Phi(U\psi U^\dagger)$, and finally Berain uncomputing the unitary U. The new protocol has the desired fidelity*

$$\begin{aligned} F_{\mathrm{wc}}(\hat{\Phi}) &= \inf_{|\psi\rangle\in\mathcal{H}_d} \int_U \langle\psi|U^\dagger\Phi(U\psi U^\dagger)U|\psi\rangle \mathrm{d}U && (3.13)\\ &= \inf_{|\psi\rangle\in\mathcal{H}_d} \int_\phi \langle\phi|\Phi(\phi)|\phi\rangle \mathrm{d}\phi && (3.14)\\ &= f(\Phi)\,. && (3.15) \end{aligned}$$

We stress again that this is only an intuitive check to justify the use of the worst-case fidelity as a figure of merit, but the above integrals over classical-to-quantum maps are not well defined.

Protocol 4 (d, N, F) Deterministic Port-Based State Preparation.

Setting. Both Anboto and Berain each have N qudits in $A_1, A_2, \ldots, A_N$ and $B_1, B_2, \ldots, B_N$ respectively, which are denoted by A and B as a whole. They share an arbitrary entangled resource in AB. There is a single round of classical communication from Anboto to Berain. Berain's correction operations are restricted to tracing out registers.

Inputs. Anboto receives the classical description of a d-dimensional quantum state.

Goal. Anboto's input state to be prepared in an arbitrary output port of Berain, with fidelity F.

The protocol:

1: Anboto performs a N outcome POVM measurement on her half of the N shared resource states, which can depend on the classical description of her state.

2: Anboto always communicates a port $x \in [N]$ to Berain with $\log(N)$ bits of classical communication.

3: Berain traces out all but register B_x.

As mentioned earlier, Anboto holding the complete classical description of the desired state allows her to construct a measurement that depends on it. In particular, this allows her to perform independent measurement operations in each port, which turns out to be enough to obtain a probability of success that exponentially increases in terms of the number of ports. Intuitively, this is equivalent to independent coin tossing, something that we are not allowed in the original port-based teleportation tasks.

Theorem 3.3.1. *There exists a (d, N, p) probabilistic PBSP protocol with probability of success*

$$p = 1 - \left(1 - \frac{1}{d}\right)^N .$$

Proof. Let us assume that Anboto and Berain share N independent maximally entangled qudits as a resource. Given a description of a d-dimensional state $|\psi\rangle$, denoted ψ, Anboto's measurement will consist of checking if the state is in any port. This can be done by just projecting to the conjugate $|\psi^*\rangle$, because for every register $i \in [N]$ we have

$$\operatorname{Tr}_{A_i}\Big[(|\psi^*\rangle\langle\psi^*|_{A_i} \otimes I_{B_i})|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\Big] = \frac{1}{d}|\psi\rangle\langle\psi|_{B_i}, \tag{3.16}$$

$$\operatorname{Tr}_{A_i}\Big[((I - |\psi^*\rangle\langle\psi^*|_{A_i}) \otimes I_{B_i})|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\Big] = \frac{1}{d}(I - |\psi\rangle\langle\psi|)_{B_i} . \tag{3.17}$$

Anboto's partial measurement can be normalized to obtain the POVM $\mathcal{M} = \{M_x(\psi)\}_{x=0}^N$ with elements

$$M_0(\psi) := \bigotimes_{i=1}^N (I - |\psi^*\rangle\langle\psi^*|)_{A_i}, \tag{3.18}$$

$$M_x(\psi) := \sum_{\substack{S\subseteq[N] \\ \text{s.t.}\, x\in S}} \left[\frac{1}{|S|}\bigotimes_{i\in S}|\psi^*\rangle\langle\psi^*|_{A_i} \otimes \bigotimes_{i\notin S}(I - |\psi^*\rangle\langle\psi^*|)_{A_i}\right], \quad \text{for } x \in [N] . \tag{3.19}$$

Outcome $x = 0$ means that the state has not been found in any port, and thus the protocol has failed. After any other outcome $x \in [N]$, Berain's system in register B_x will be in the desired state

$$\begin{aligned}
\rho_{B_x} &= \frac{\operatorname{Tr}_{AB_{\bar{x}}}\Big[M_x(\psi) \bigotimes_{i=1}^N |\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\Big]}{\operatorname{Tr}\Big[M_x(\psi) \bigotimes_{i=1}^N |\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\Big]} && (3.20) \\
&= d \operatorname{Tr}_{A_x}\Big[(|\psi^*\rangle\langle\psi^*|_{A_x} \otimes I_{B_x})|\phi_d^+\rangle\langle\phi_d^+|_{A_xB_x}\Big] && (3.21) \\
&= |\psi\rangle\langle\psi|_{B_x} . && (3.22)
\end{aligned}$$

Given this measurement, the probability of success of the protocol can be calculated easily as we have a closed expression for the failing measurement

$$p = 1 - \mathrm{Tr}\left[M_0(\psi) \bigotimes_{i=1}^{N} |\phi_d^+\rangle\langle\phi_d^+|_{A_i B_i} \right] = 1 - \left(1 - \frac{1}{d}\right)^N . \tag{3.23}$$

□

We can always construct a deterministic PBSP protocol from a probabilistic PBSP by just outputting a random port when a failure is measured. Although in the PBT protocol this construction gives us a greater fidelity than the probability of success, see eq. (3.10), for the PBSP case this is not necessarily true. For example, in the measurement we described above the extra bit is always orthogonal to the desired state, thus the deterministic PBSP obtained from spreading the failure evenly among the ports leads to the same channel fidelity as the probability of success of the probabilistic protocol described in Theorem 3.3.1.

Theorem 3.3.2. *There exists a (d, N, F) deterministic PBSP protocol with channel fidelity*

$$F_{\mathrm{wc}} = 1 - \left(1 - \frac{1}{d}\right)^N .$$

Proof. Let us assume that Anboto and Berain share N independent maximally entangled qudits as a resource. Given a description of a d-dimensional state $|\psi\rangle$, denoted ψ, Anboto's measurement will consist of checking if the state is in any port via the measurements

$$M_x'(\psi) = M_x(\psi) + \frac{1}{N} \bigotimes_{i=1}^{N} (I - |\psi^*\rangle\langle\psi^*|)_{A_i}, \quad \text{for } x \in [N] , \tag{3.24}$$

where the $M_x(\psi)$'s are the same as in eq. (3.18). Any of the outcomes $x \in [N]$ will now occur with equal probability

$$\begin{aligned} p_x &= \mathrm{Tr}\left[(M_x'(\psi) \otimes I_B) \bigotimes_{i=1}^{N} |\phi_d^+\rangle\langle\phi_d^+|_{A_i B_i} \right] && (3.25) \\ &= \frac{1}{d} \sum_{i=0}^{N-1} \frac{\binom{N-1}{i}}{i+1} \left(\frac{1}{d}\right)^i \left(1 - \frac{1}{d}\right)^{(N-1)-i} + \frac{1}{N}\left(1 - \frac{1}{d}\right)^N , && (3.26) \end{aligned}$$

and after outcome $x \in [N]$ the state in Berain's register B_x will be

$$\rho_{B_x} = \frac{1}{p_x}\left(\frac{1}{d}|\psi\rangle\langle\psi|_{B_x}\sum_{i=0}^{N-1}\frac{\binom{N-1}{i}}{i+1}\left(\frac{1}{d}\right)^i\left(1-\frac{1}{d}\right)^{(N-1)-i} + \frac{1}{Nd}(I-|\psi\rangle\langle\psi|)_{B_x}\left(1-\frac{1}{d}\right)^{N-1}\right). \quad (3.27)$$

From the above expression we can see that for every input ψ the fidelity between the outcome of the protocol $\Phi(\psi)$ and the desired state $|\psi\rangle$ will be equivalent to the probability of success p, as the latest term is orthogonal to $|\psi\rangle_{B_x}$ and thus will not contribute to the overlap between the states. Formally,

$$\begin{aligned} F(\Phi(\psi), |\psi\rangle) &= \sum_{x=1}^{N} p_x \langle\psi|_{B_x}\rho_{B_x}|\psi\rangle_{B_x} && (3.28)\\ &= \frac{N}{d}\sum_{i=0}^{N-1}\frac{\binom{N-1}{i}}{i+1}\left(\frac{1}{d}\right)^i\left(1-\frac{1}{d}\right)^{(N-1)-i} && (3.29)\\ &= 1-\left(1-\frac{1}{d}\right)^N. && (3.30)\end{aligned}$$

□

Note that the accuracy of PBSP scales exponentially in terms of the number of ports, both for the probabilistic and the deterministic case. This contrasts the PBT case; where the upper bound for both terms is linear in terms of number of ports. In other words, having the classical description of the state gives the sender enough power to go from linear to exponential accuracy in the number of ports; or that PBSP is a strictly easier task to perform than PBT.

3.4 Universal Programmable Hybrid Processors

Quantum computers have the ability to perform arbitrary unitary operations in two-level systems, i.e., we can decompose any unitary acting on qubits in gate arrays which can be implemented using finite resources. Moreover, classical computers can perform arbitrary operations on any input data with a fixed universal gate array, which can be programmed to perform any operation on the input data. However, this was shown to be impossible quantumly by Nielsen and Chuang [NC97]: the so-called *no-programming theorem* states that perfect universal quantum processors require an orthogonal state for each unitary that we desire to perform, and since there are infinitely many unitaries acting on any d-level system, we would need infinite-dimensional spaces to perform arbitrary unitaries on quantum states. The relaxed scenario of approximate UPQPs has also attracted much attention [HZB06, KPPG19, Maj18], but a recent article by Yang, Chiribella and Renner [YRC20] closed the gap between the upper and lower bound on the memory dimension m of a d-dimensional ε-approximate UPQP:

$$\left(\frac{1}{\varepsilon}\right)^{\frac{d^2-1}{2}} \geq m \geq \left(\frac{1}{\varepsilon}\right)^{\frac{d^2-1}{2}-\alpha}, \quad \text{for all } \alpha > 0\,. \tag{3.31}$$

In other words, the memory dimension scales with the dimension-square exponentially in the error, which nearly saturates the upper bound in [KPPG19] given by the ε-net of the unitaries. Intuitively, it is as saying that the memory cannot do much better than pointing to a unitary from the ε-net, which is in a way a classical operation.

Since for some applications the user of the quantum program might actually know which state they would like to apply the (to them unknown) unitary to, we imagine an analogous task to ε-UPQPs, where we want to program a universal gate array but to which we will give a classical input and expect to perform an arbitrary unitary operation on the quantum state the classical input specifies. The initial state of our classical-quantum system will be of the form

$$(\psi, |\phi\rangle_M)\,, \tag{3.32}$$

where ψ will refer to the classical description of the d-dimensional quantum state $|\psi\rangle$ we want to perform the operation on, w.l.o.g. we assume it to be pure, and $|\phi\rangle_M$ is the state of the m-dimensional program register. For any input, the processor P will then map this system to

$$(\psi, |\phi\rangle_M) \mapsto \mathsf{P}(\psi, |\phi\rangle_M)\,, \tag{3.33}$$

and we say that the processor P implements a unitary U in the data register $|\psi\rangle_D$ specified by the classical input ψ if

$$\mathrm{Tr}_M[\mathsf{P}(\psi, |\phi\rangle_M)] = U|\psi\rangle\langle\psi|_D U^\dagger\,. \tag{3.34}$$

Note that although the processor itself is not unitary, since it takes classical and quantum data, we want eq. (3.34) to hold for *any input state*, thus it is not possible to just encode all the information about the unitary desired in the classical data. Let us denote by $\mathsf{P}_\psi := \mathsf{P}(\psi, \cdot)$ the universal gate array specified by classical input ψ.

As in the case of ε-UPQPs, we are interested in an approximate version of this processor. We will call such a system a *approximate universal programmable hybrid processor* (ε-UPHP) to refer to the classical-input quantum-output dynamics of the system explicitly. We can also distinguish between probabilistic and deterministic ε-UPHPs, in analogy to the different variants of port-based teleportation, but since every probabilistic processor can be converted into a deterministic one, it is more common to treat the latter. We will do the same here for the sake of simplicity.

Definition 3.4.1 (ε-UPHP)**.** *A family of quantum channels* $\mathsf{P} := \{P_\psi\}_\psi \subset \mathrm{CPTP}(\mathcal{H}_M, \mathcal{H}_D \otimes \mathcal{H}_M)$ *is a d-*dimensional ε-Universal Programmable Hybrid Processor*, if for every unitary $U \in \mathrm{U}(\mathcal{H}_D)$ there exists a unit vector $|\phi_U\rangle \in \mathrm{D}(\mathcal{H}_M)$ such that*

$$\frac{1}{2}\left\| \mathrm{Tr}_M[\mathsf{P}_\psi(|\phi_U\rangle\langle\phi_U|_M)] - U|\psi\rangle\langle\psi|_D U^\dagger \right\|_1 \leq \varepsilon \quad \textit{for all input } \psi\,,$$

where ψ is the classical description of the d-dimensional quantum state $|\psi\rangle_D \in \mathrm{D}(\mathcal{H}_D)$.

We describe the above accuracy parameter ε in terms of the trace norm, equivalent to the average fidelity, because this one is easier to work with, and we want to draw a parallel with the accuracy parameter in the diamond norm for the original UPQP. However, we believe that the symmetric structure of this task analogous to PBT will allow us to show that the optimal accuracy in terms of diamond norm will be equivalent to the one worst-case fidelity and average fidelity as in [YRC20].

We now give an existence result by constructing a ε-UPHP from PBSP. The UPHP constructed has a lower scaling of the memory dimension m in terms of the error ε and the data dimension d, than the optimal scaling for ε-UPQPs. We obtain this by lifting the exponential properties of PBSP, which once again shows that the former classical-quantum task is easier to perform by the processor than the fully quantum one, because we are not constrained by the no-programming theorem anymore.

Theorem 3.4.2. *There exists a d-dimensional ε-UPHP, $\mathsf{P} := \{\mathsf{P}_\psi\}_\psi \subset \mathrm{CPTP}(\mathcal{H}_M, \mathcal{H}_D \otimes \mathcal{H}_M)$, with memory dimension*

$$m \leq \left(\frac{1}{\varepsilon}\right)^{4d\ln(d)}.$$

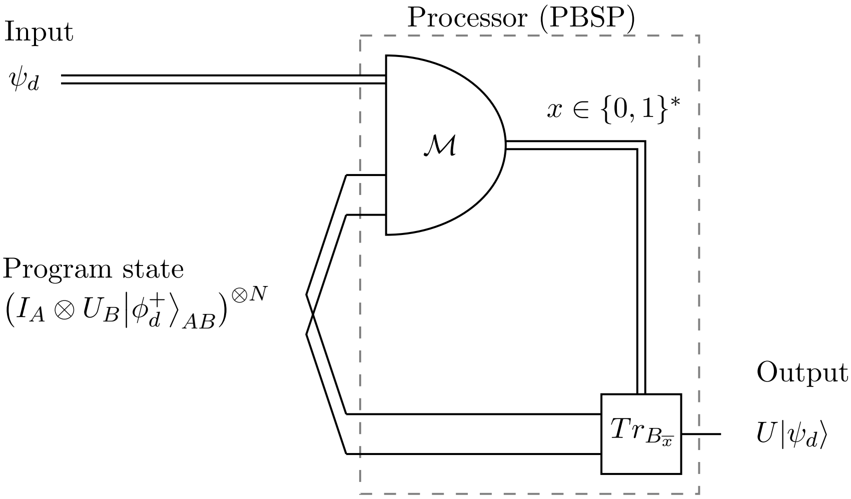

Figure 3.2: Construction of approximate UPHP from PBSP.

Proof. Given a unitary $U \in \mathrm{U}(\mathcal{H}_D)$, we can adapt the PBSP protocol given in Theorem 3.3.2 to transmit the state with the unitary applied to all ports instead, so that the ε-UPHP will be a machine simulating the PBSP interaction between Anboto and Berain, where Anboto's classical PBSP input will be the one given to the processor, and the PBSP resource state shared between Anboto and Berain will be used as the program state of the processor. Therefore, by using the Choi state of the unitary as a memory state, i.e. by using a PBSP protocol with N copies of the state $(I_A \otimes U_B)|\phi_d^+\rangle_{AB}$ as a resource state, the measurements described by Theorem 3.3.2 will output the unitary applied to the input state $U|\psi\rangle\langle\psi|U^\dagger$ with high fidelity.

By the properties of the maximally entangled state, we have that by projecting the unitary U applied to the maximally entangled state to the conjugate of the desired state $|\psi^*\rangle$, we obtain the unitary applied to the state with the same probability as projecting without the unitary. This holds because

$$\begin{aligned}\mathrm{Tr}_{A_i}\Big[(|\psi^*\rangle\langle\psi^*|_{A_i} \otimes I_{B_i})(I_{A_i} \otimes U_{B_i})|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}(I_{A_i} \otimes U_{B_i}^\dagger)\Big]& \\ = \frac{1}{d}U|\psi\rangle\langle\psi|_{B_i}U^\dagger\,, & \quad (3.35)\end{aligned}$$

$$\mathrm{Tr}_{A_i}\Big[((I - |\psi^*\rangle\langle\psi^*|_{A_i}) \otimes I_{B_i})(I_{A_i} \otimes U_{B_i})|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}(I_{A_i} \otimes U_{B_i}^\dagger)\Big]$$

$$= \frac{1}{d} U(I - |\psi\rangle\langle\psi|)_{B_i} U^\dagger \,. \quad (3.36)$$

Therefore, when Anboto makes the measurements $\{M_x(\psi)\}_{x=1}^N$ as described in eq. (3.18), the probability of success of port $x \in [N]$ will be the same as the original PBSP eq. (3.25), but the outcome state will now have the unitary applied to it

$$\rho_{B_x} = \frac{1}{p_x}\left(\frac{1}{d} U|\psi\rangle\langle\psi|_{B_x} U^\dagger \sum_{i=0}^{N-1} \frac{\binom{N-1}{i}}{i+1}\left(\frac{1}{d}\right)^i \left(1-\frac{1}{d}\right)^{(N-1)-i} \right.$$
$$\left. + \frac{1}{Nd} U(I - |\psi\rangle\langle\psi|)_{B_x} U^\dagger \left(1-\frac{1}{d}\right)^{N-1}\right) . \quad (3.37)$$

Finally, the outcome fidelity between the protocol $\Phi_U(\psi)$ and the unitary applied to the system $U|\psi\rangle$ will be

$$F(\Phi_U(\psi), U|\psi\rangle) = \sum_{x=1}^N p_x \langle\psi|_{B_x} U^\dagger \rho_{B_x} U|\psi\rangle_{B_x} \quad (3.38)$$

$$= \frac{N}{d}\sum_{i=0}^{N-1} \frac{\binom{N-1}{i}}{i+1}\left(\frac{1}{d}\right)^i \left(1-\frac{1}{d}\right)^{(N-1)-i} \quad (3.39)$$

$$= 1 - \left(1-\frac{1}{d}\right)^N . \quad (3.40)$$

We have thus constructed a processor P such that for every unitary $U \in \mathrm{U}(\mathcal{H}_D)$, there exists a memory state

$$|\phi_U\rangle_M := (I_A \otimes U_B) \bigotimes_{i=1}^N |\phi_d^+\rangle_{A_i B_i} \,, \quad (3.41)$$

for which

$$F(\mathrm{Tr}_M[\mathsf{P}_\psi(|\phi_U\rangle_M)], U|\psi\rangle_D) = F(\Phi_U(\psi), U|\psi\rangle) = 1 - \left(1-\frac{1}{d}\right)^N , \quad (3.42)$$

for all inputs ψ. The dimension of the memory m of the ε-UPHP is the dimension of the maximally mixed state d^2 to the power of the number of ports N needed $m := d^{2N}$. Therefore, if the number of ports is chosen such that $N \geq d\ln\left(1/\varepsilon^2\right)$, from the properties of the exponential function and the Fuchs–van de Graaf inequality eq. (3.7), we obtain

$$\frac{1}{2}\left\|\Phi_U(\psi) - U|\psi\rangle\langle\psi|U^\dagger\right\|_1 \leq \sqrt{1 - F(\Phi_U(\psi), U|\psi\rangle)} \tag{3.43}$$

$$\leq \left(1 - \frac{1}{d}\right)^{d\frac{\ln\left(1/\varepsilon^2\right)}{2}} \leq e^{\frac{\ln\left(\varepsilon^2\right)}{2}} = \varepsilon\,, \tag{3.44}$$

which means that in order to obtain a processor with accuracy ε, it is enough for the memory to be of dimension

$$m \leq \left(\frac{1}{\varepsilon}\right)^{4d\ln(d)}. \tag{3.45}$$

□

3.5 Achievable optimality of the tasks

The existence results of the previous chapters show that having a classical description of a quantum state substantially increases the power of teleportation and processing tasks. In this section, we complement these with bounds on the optimality achievable by port-based state preparation (PBSP) and universal programmable hybrid processors (UPHP).

3.5.1 Approximate UPHP

As we saw previously, an approximate UPHP is a quantum machine that, given a description of a quantum state, outputs a quantum state to which a unitary of choice is applied. In a way, ε-UPHPs are encoding the functionality of the unitary in a quantum memory state, but this should intuitively also involve encoding at least some classical information into this quantum state.

We formalize this intuition in the following theorem, where we show the hardness of ε-UPHP though the hardness of quantum random access codes (QRACs). That is, we show that UPHPs can be used as QRACs, which are limited by Nayak's bound, Theorem 3.2.5.

Theorem 3.5.1. *For any d-dimensional ε-UPHP, that is $\mathsf{P} := \{\mathsf{P}_\psi\}_\psi \subset$ $\mathrm{CPTP}(\mathcal{H}_M, \mathcal{H}_D \otimes \mathcal{H}_M)$, the memory dimension must be at least*

$$m \geq 2^{\frac{d}{2}(1-h(2\varepsilon))}\,.$$

Proof. Given a (d, m, ε)-UPHP we can construct a $(d/2, \log(m), 1-2\varepsilon)$-QRAC. Let us consider an arbitrary $d/2$ bit-string $f \in \{0,1\}^{d/2}$ where each bit $f(x)$ is indexed by the bit-string $x \in \{0,1\}^{\log(d)-1}$. we can equivalently see this as a description of a Boolean function $f : \{0,1\}^{\log(d)-1} \rightarrow \{0,1\}$.

We can compress this Boolean function into a $\log(m)$-dimension Hilbert space using the (d, m, ε)-UPHP. That is, note that every Boolean function f can be computed by a unitary

$$U_f \colon |x\rangle_1 \otimes |y\rangle_2 \mapsto |x\rangle_1 \otimes |y \oplus f(x)\rangle_2, \quad \text{with } y \in \{0,1\}\,, \tag{3.46}$$

where the unitary associated with the Boolean functions of input size $\log(d)-1$ can be hosted in dimension d.

Now from the ε-UPHP assumption in eq. (3.34), given $U_f \in \mathrm{U}(\mathcal{H}^{d/2})$ there exists a memory state $|\phi_f\rangle \in \mathcal{H}_M$ such that the processor $\mathsf{P} := \{\mathsf{P}_\psi\}_\psi \subset \mathrm{CPTP}(\mathcal{H}_M, \mathcal{H}_D \otimes \mathcal{H}_M)$ approximates the desired unitary

$$\frac{1}{2}\left\|\mathrm{Tr}_M\left[\mathsf{P}_x(|\phi_f\rangle\langle\phi_f|_M)\right] - U_f(|x\rangle\langle x| \otimes |0\rangle\langle 0|)U_f^\dagger\right\|_1 \leq \varepsilon\,. \tag{3.47}$$

Now consider the channel consisting of applying P_x, tracing out the parts we do not need, and then applying a measurement in the standard basis. Via a standard argument, we can write this procedure as a single measurement which we call M_x^0 and M_x^1. The exact form can be derived e.g. through writing out the Kraus operators. In particular, if define K_k^x to be the Kraus operators of the channel P and subsequent partial trace, we have that $\mathrm{Tr}_{1,M}[\mathsf{P}_x(\cdot)] = \sum_k K_k^x \rho (K_k^x)^\dagger$. Now define the set of $d/2$ measurements $\{M_x^0, M_x^1\}_{x \in \{0,1\}^{\log(d)-1}}$:

$$M_x^y := \sum_k (K_k^x)^\dagger |y\rangle\langle y| K_k^x, \quad \text{for } y \in \{0,1\}\,. \tag{3.48}$$

We can verify that for these measurements it holds that

$$\begin{aligned}
\mathrm{Tr}[M_x^y \rho] = \mathrm{Tr}\left[\sum_k (K_k^x)^\dagger |y\rangle\langle y| K_k^x \rho\right] &= \mathrm{Tr}\left[|y\rangle\langle y| \sum_k K_k^x \rho (K_k^x)^\dagger\right] &(3.49)\\
&= \mathrm{Tr}[|y\rangle\langle y|\, \mathrm{Tr}_{1,M}[\mathsf{P}_x(\rho)]]\,, &(3.50)
\end{aligned}$$

as desired.

We can use these as the set of measurements of a $(d/2, \log(m), 1-2\varepsilon)$-QRAC, because the probability of correctly guessing the bit $f(x)$ of the bit-string f is at least

$$
\begin{aligned}
&\operatorname{Tr}\left[M_x^{f(x)}|\phi_f\rangle\langle\phi_f|\right] && (3.51)\\
&= \operatorname{Tr}[|f(x)\rangle\langle f(x)| - |f(x)\rangle\langle f(x)|] + \operatorname{Tr}\left[M_x^{f(x)}|\phi_f\rangle\langle\phi_f|\right] && (3.52)\\
&= 1 - \operatorname{Tr}[|f(x)\rangle\langle f(x)| - \operatorname{Tr}_1[(I_1 \otimes |f(x)\rangle\langle f(x)|_2) \operatorname{Tr}_M[\mathsf{P}_x(|\phi_f\rangle\langle\phi_f|)]]] && (3.53)\\
&\geq 1 - \| |f(x)\rangle\langle f(x)| - \operatorname{Tr}_1[(I_1 \otimes |f(x)\rangle\langle f(x)|_2) \operatorname{Tr}_M[\mathsf{P}_x(|\phi_f\rangle\langle\phi_f|)]] \|_1 && (3.54)\\
&\geq 1 - 2\varepsilon\,, && (3.55)
\end{aligned}
$$

where in the last line we used the data processing inequality for the trace distance.

Finally, from Nayak's bound, Theorem 3.2.5, we know that, even when allowing some error, we cannot compress bits into less qubits, which leads to the desired bound

$$\log(m) \geq \frac{d}{2}(1 - h(2\varepsilon))\,, \tag{3.56}$$

where h refers to the binary entropy function. □

We remark that this bound (which of course also applies to ε-UPQP, besides ε-UPHP) is tighter than the bound in [KPPG19], that was proven in the context of ε-UPQPs. We also want to point out that we were not able to reproduce the impossibility results blowing-up with the error tending to zero present in the UPQP literature [HZB06, Maj18, PG06]. These results seem to rely on the spectral decomposition of the processor, but the fact that we now have infinite copies of the input state makes it impossible to exploit the uniqueness of the spectral decomposition. In fact, it is not even clear to us how to prove the no-programming theorem as in [NC97] for the zero-error UPHP case.

3.5.2 PBSP

In the following, we give two bounds on the limitations of Port-Based State Preparation. The first bound follows from the achievable optimality of universal programmable hybrid processors (UPHP), recall that the construction of PBSP from Theorem 3.4.2 relied on UPHPs. The second bound follows from a non-signalling argument.

Theorem 3.5.2. *The optimal fidelity of a (d, N, F^*)-Deterministic Port-Based State Preparation protocol with N independent resource states is upper*

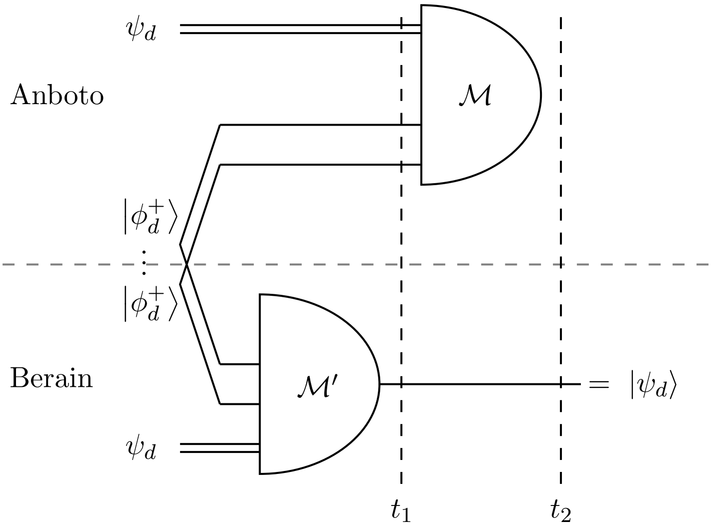

Figure 3.3: Non-signalling times t_1 before Anboto's measurement and t_2 after Anboto's measurement.

bounded by

$$1 - h(\sqrt{1 - F^*}) \leq \frac{4N \log(d)}{d} \,.$$

Proof. Recall that we saw in Theorem 3.4.2 how to use a $(d, N, F = 1 - \varepsilon^2)$-PBSP to construct a d-dimensional ε-UPHP with memory dimension $m = d^{2N}$. Moreover, by Theorem 3.5.1 the dimension of the program register of d-dimensional ε-UPHP is lower bounded by

$$d^{2N} \geq 2^{\frac{d}{2}(1-h(2\varepsilon))} \,. \tag{3.57}$$

Substituting $\varepsilon = (\sqrt{1 - F})/2$, we obtain the desired bound

$$1 - h\Big(\sqrt{1 - F}\Big) \leq \frac{4N \log(d)}{d} \,. \tag{3.58}$$

□

For the non-signaling argument, the idea is for Berain to make a measurement on his half of the shared resources, and to compare the outcome before and after Anboto makes a measurement, without any communication. Note that the following bounds explicitly use the fact that the resource states are maximally entangled states.

Theorem 3.5.3. *The optimal probability of a (d, N, p^*)-Probabilistic Port-Based State Preparation protocol with N maximally entangled resource states is upper bounded by*

$$p^*_{\mathrm{EPR}} \leq 1 - \left(1 - \frac{1}{d}\right)^N .$$

Proof. Imagine that Anboto wishes to execute a PBSP protocol to transport an arbitrary d-dimensional state $|\psi\rangle$, that for the sake of the argument is also known to Berain.

Imagine that Berain measures all his halfs of the N maximally mixed states to see if they are in the state $|\psi\rangle\langle\psi|$ or $I - |\psi\rangle\langle\psi|$, for example by using the POVM $\mathcal{M}' = \{M_\top, M_\perp\}$, with

$$M_\perp(\psi) = \bigotimes_{i=1}^{N} (I - |\psi\rangle\langle\psi|)_{B_i} \quad \text{and} \quad M_\top = I - M_\perp \,. \tag{3.59}$$

We will compare Berain's outcome before and after Anboto does any measurement. Before Anboto does any measurement, see time t_1 in figure 3.3, observe because $|\psi\rangle$ is random, any specific port will have probability $\frac{1}{d}$ to see $|\psi\rangle$. Therefore, because the resource state is a product state of identical ports, the probability of seeing $|\psi\rangle$ in *some* port is

$$\begin{aligned} p^{t_1}_{\text{success}} &= 1 - \mathrm{Tr}\left[(I_A \otimes M_\perp) \bigotimes_{i=1}^{N} |\phi^+_d\rangle\langle\phi^+_d|_{A_i B_i} \right] & (3.60) \\ &= 1 - \prod_{i=1}^{N} \mathrm{Tr}\Big[(I_{A_i} \otimes (I - |\psi\rangle\langle\psi|)_{B_i}) |\phi^+_d\rangle\langle\phi^+_d|_{A_i B_i} \Big] & (3.61) \\ &= 1 - \left(1 - \frac{1}{d}\right)^N . & (3.62) \end{aligned}$$

Let $\mathcal{M} = \{M_0, M_1, \ldots, M_N\}$ be Anboto's POVM, outcome $x \in [N]$ refers to the state being in port B_x, and outcome $x = 0$ refers to the protocol failing. After Anboto does her measurement, see time t_2 in figure 3.3, if the probabilistic PBSP went correctly, a port B_x would exist containing $|\psi\rangle$. Therefore, the probability that Berain sees any port with the teleported state has to be at least p^*_{EPR} by correctness of the protocol. Formally, the probability of Berain seeing $|\psi\rangle$ in *some* port is

$$p^{t_2}_{\text{success}} = \operatorname{Tr}\left[(I_A \otimes M_\top) \bigotimes_{i=1}^{N} |\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right] \tag{3.63}$$

$$\begin{aligned} &= \operatorname{Tr}\left[(M_0 \otimes M_\top) \bigotimes_{i=1}^{N} |\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right] \\ &\quad + \operatorname{Tr}\left[\left(\sum_{x=1}^{N} M_x \otimes M_\top\right) \bigotimes_{i=1}^{N} |\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right] \end{aligned} \tag{3.64}$$

$$= \operatorname{Tr}\left[(M_0 \otimes M_\top) \bigotimes_{i=1}^{N} |\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right] + p^*_{\text{EPR}}\,. \tag{3.65}$$

By non-signalling, the probabilities of some event happening on Berain's side (if no communication is involved) is identical before and after Anboto's measurement. So by equating the probabilities of the event at time t_1 and t_2 we obtain the desired bound

$$1 - \left(1 - \frac{1}{d}\right)^N = p^{t_1}_{\text{success}} = p^{t_2}_{\text{success}} \geq p^*_{\text{EPR}}\,. \tag{3.66}$$

□

The same argument can be applied for deterministic PBSP, where the analogous non-signalling theorem for fidelity tells us that the entanglement fidelity of a protocol cannot increase if no communication happens between the parties. We will only include a sketch of the proof.

Theorem 3.5.4. *The optimal fidelity of a (d, N, F^*)-Deterministic Port-Based State Preparation protocol with N maximally entangled resource states is upper bounded by*

$$F^*_{\text{EPR}} \leq 1 - \left(1 - \frac{1}{d}\right)^N .$$

Proof. We will compare Berain's outcome when performing the measurement eq. (3.59) before and after Anboto does the optimal measurement. Before Anboto performs any measurement, see time t_1 in figure 3.3, the probability of seeing $|\psi\rangle$ in some port is

$$p^{t_1}_{\text{success}} = 1 - \left(1 - \frac{1}{d}\right)^N . \tag{3.67}$$

We can bound the probability of projecting into a tensor product $\bigotimes_{i=1}^{N}(I - |\psi\rangle\langle\psi|)_{B_i}$ by the minimum of all the marginal projections, which gives us an upper bound for the probability of failure in terms of individual registers

$$p_{\text{failure}}^{t_1} = 1 - p_{\text{success}}^{t_1} \tag{3.68}$$

$$= \mathrm{Tr}\left[\left(I_A \otimes \bigotimes_{i=1}^{N}(I - |\psi\rangle\langle\psi|)_{B_i}\right)\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right] \tag{3.69}$$

$$\leq \min_{x\in[N]} \mathrm{Tr}\left[\left(I_{AB_{\bar{x}}} \otimes (I - |\psi\rangle\langle\psi|)_{B_x}\right)\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right] \tag{3.70}$$

$$= \min_{x\in[N]} (I - |\psi\rangle\langle\psi|)_{B_x} \mathrm{Tr}_{AB_{\bar{x}}}\left[\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right]. \tag{3.71}$$

Let $\mathcal{M} = \{M_1, \dots, M_N\}$ be Anboto's POVM, outcome $x \in [N]$ refers to the state being in port B_x. By non-signalling a measurement on Anboto's side without communicating the outcomes to Berain cannot change his perspective of the state, this is, for every $x \in [N]$:

$$\sum_{i\in[N]} \mathrm{Tr}_{AB_{\bar{x}}}\left[(M_i(\psi) \otimes I_B)\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right] = \mathrm{Tr}_{AB_{\bar{x}}}\left[\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right]. \tag{3.72}$$

Therefore we can rewrite the upper bound on the probability of failure as

$$p_{\text{failure}}^{t_1} = \min_{x\in[N]} \sum_{i\in[N]} (I - |\psi\rangle\langle\psi|)_{B_x} \mathrm{Tr}_{AB_{\bar{x}}}\left[(M_i(\psi) \otimes I_B)\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right] \tag{3.73}$$

$$\leq \sum_{x\in[N]} (I - |\psi\rangle\langle\psi|)_{B_x} \mathrm{Tr}_{AB_{\bar{x}}}\left[(M_x(\psi) \otimes I_B)\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right]. \tag{3.74}$$

Note that by the correctness of the protocol, this value is related to the fidelity in the desired way

$$1 - F_{\text{EPR}}^{*} \tag{3.75}$$

$$= 1 - \sum_{x\in[N]} \langle\psi|_{B_x} \mathrm{Tr}_{AB_{\bar{x}}}\left[(M_x(\psi) \otimes I_B)\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right]|\psi\rangle_{B_x} \tag{3.76}$$

$$= \sum_{x\in[N]} p_x - \mathrm{Tr}\left[|\psi\rangle\langle\psi|_{B_x} \mathrm{Tr}_{AB_{\bar{x}}}\left[(M_x(\psi)\otimes I_B)\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right]\right] \tag{3.77}$$

$$= \sum_{x\in[N]} \mathrm{Tr}\left[(I-|\psi\rangle\langle\psi|)_{B_x} \mathrm{Tr}_{AB_{\bar{x}}}\left[(M_x(\psi)\otimes I_B)\bigotimes_{i=1}^{N}|\phi_d^+\rangle\langle\phi_d^+|_{A_iB_i}\right]\right]. \tag{3.78}$$

□

3.6 Conclusion

In this chapter we initiate the study of known-state variants of port-based teleportation and universal programmable processors. We show that if Anboto knows the state she intends to send, while Berain can only perform classical computations, then it is exponentially better in the error dependence to use this knowledge than to first prepare the state and then perform standard PBT. We later show similar results for universal programmable processors, which can also make much better use of their memory resources if they have a classical description of the state on which they have to perform operations.

These results do not come as a surprise, as it is common for quantum information to be much more difficult to access than classical information. However, the tasks introduced in this chapter provide a natural setting to compare resource trade-offs.

Since PBT has attracted so much attention and found so many applications, it is natural to ask which of these translate to the known-state variant, and if they become easier to perform. In particular, some of these questions might be:

- Exact formulas for the fully optimized measurement and resource exist for both probabilistic and deterministic PBT, derived using graphical algebra [SSMH17, MSSH18], and asymptotic term expressions derived from these are discussed in [CLM$^+$21, PG13, MSSH18]. Here we only calculate the optimal achievable error of PBSP when maximally entangled states are used as a resource, but can such an analysis for the fully-optimized case be performed for PBSP?

- Recently some efficient algorithms for computing PBT were presented which exploit the symmetry of the task [GBO23, FTH23]. Can we also exploit the symmetry of PBSP to provide efficient algorithms for its implementation?

- Note that we do not close the gap between the upper and lower bound in the memory dimension for UPHP. In fact, even though we produce scaling in the dimension which is close to tight, we were unable to reproduce the known lower-bound results for UPQP which show that the dimension blows up as the error ε tends to zero, which should intuitively hold for UPHP as well.

- From the perspective of high-order physics, it turns out that for probabilistic exact unitary transformations, [QDS$^+$19a] showed that the maximal success probability can always be obtained by probabilistic PBT, whilst this is not the case for the deterministic case. What happens in the case of a known state?

Chapter 4

The Threshold Authentication Scheme

This chapter is based on the article *An efficient combination of quantum error correction and authentication* [DMS25], and is joint work with Yfke Dulek and Florian Speelman.

The preliminaries have been shortened as the security model has already been treated in Section 2.2. The remaining sections have been edited for style and typographical consistency.

In this chapter we study the possibility of finding quantum authentication protocols that are inherently robust against noise. In Section 4.3 we study the efficiency, in terms of number of qubits required, for the existing method of combining an authentication code with an error-correcting code. In Section 4.4 we introduce the *threshold* authentication scheme, prove its correctness and security, and show that for a certain range of parameters it is indeed more efficient than the previous method.

4.1 Introduction

Authentication is one of the most fundamental tasks of modern cryptography — for many applications it is imperative that the integrity of data is preserved, not just against noise and random errors, but even against adversarial attacks. Constructions for message authentication codes (MACs) underlay many important cryptographic protocols that are in constant use for secure internet communication. We study the notion of *quantum authentication* introduced by Barnum, Crepeau, Gottesman, Smith, and Tapp [BCG$^+$02], where instead of wanting to ascertain the integrity of classical data, the data involved consists of qubits.

Authentication is usually applied to messages that will be transmitted at some point, and such a transmission involves incurring some *error* by the quantum channel which is used. Although the MACs present in the literature will inevitably reject whenever any error is present in the channel, it is possible to first encode the data in a quantum authentication code, and then wrap the result in an *error-correction code*, as pointed out by Hayden, Leung and Mayers [HLM16] and later proven to be secure by Portmann [Por17].

The primitives of quantum authentication and error correction have a conceptual overlap, in the sense that both aim to protect data against modifications. However, in practice there is a large difference in how they are built: authentication codes need to protect against any adversarial attack, and therefore often are extremely sensitive against even minor attempted modifications. On the other hand, an error-correcting code should be robust against typical (usually low-weight) modifications of the encoded data. Given that the goals of these codes are similar, one might wonder whether this is doing 'double work' in some sense, making the resulting encoded state larger than necessary. In this chapter, we will study the following question:

> *Is there a more efficient way of authenticating data over a noisy channel than concatenating an authentication and an error-correction code?*

We stress that the comparison is only interesting when the outer error-correcting code is indeed doing 'double encoding', as is the case for any code that encodes a single qubit of data like concatenated codes, but does not holds for 'good' error-correcting codes with linear rate and distance.

Our contributions. In this chapter, we give evidence that this is indeed the case: We construct a code which functions both as a quantum error-correcting code and as a quantum authentication code, and which is more efficient than the naive concatenation of these functionalities would imply. We call our code the *threshold code*, alluding to the fact that it is a variant of the *trap authentication code* where a threshold amount of traps is allowed to be triggered while still accepting the decoded message. In particular, this code preserves several of the useful computational properties of the original trap code, if a CSS code is used as the underlying error-correcting code.

In order to assess the performance of the code, we study its *correctness*– protection against noise– and *security* –protection against adversarial attacks– separately, and we show that the threshold code is correct and composably secure by proving that the resulting code is a good purity testing code. We do so in the abstract cryptography (AC) framework, so that the same security proof will also imply security under most other security definitions (if these do not require extra properties such as key recycling). In particular, in order to obtain ε-correctness and δ-security, there exists a threshold parameter such that the total number of qubits required to encode one qubit of data scales as

$$\Omega\left(\max\left\{\log(1/\varepsilon)^{1/\alpha}, \log(1/\delta)^{1/\alpha}\right\}\right), \tag{4.1}$$

where α is the order of decay of the error-correcting code, which is $\alpha = \log_n((d+1)/2)$ for a concatenated $[[n,1,d]]$-QECC.

Relation to the literature. Starting with the work of Barnum, Crepeau, Gottesman, Smith, and Tapp [BCG+02], several quantum authentication codes have been proposed. In this chapter, when comparing the performance of the threshold code to the naive concatenation of the authentication and error correction primitives separately, we mostly work with two prominent examples; namely the *Clifford code* and the *trap code*, not going into depth for other examples such as the polynomial code [BCG+06] or the Auth-QFT-Auth scheme [GYZ17].

The Clifford code [ABOE08] constructs a very effective authentication scheme, which involves attaching a number of flag qubits to the plaintext, and then scrambling the state using a random Clifford – this turns out to be a very efficient way of guaranteeing security, and it can also be used as a building block for interactive proofs [ABOE08] and multi-party computation [DNS10, DNS12, DGJ+20]. The trap code [BGS13, BW16] constructs a scheme, for which encoding consists of interspersing the plaintext (in an error-correcting code) with so-called traps which try to detect an adversary's attempted modifications. A very interesting property of this authentication scheme is its natural interaction with computation; it is possible to perform some quantum gates 'transversally' on qubits of the ciphertext, which results in a valid authentication of a new ciphertext (with an updated key). This property enabled the trap code to be a crucial ingredient in various results within quantum cryptography, such as the construction of quantum one-time programs [BGS13], a scheme for quantum zero-knowledge proofs for QMA [BJSW16], and verifiable homomorphic encryption [ADSS17]. Also see an extended version of the trap code which supports key recycling and ciphertext authentication [DS18] for more context of this code.

Multiple works have followed the first notions of security for the primitive of quantum authentication of Barnum et al. [BCG$^+$02], which did not consider adversaries entangled with the encrypted message. An important requirement for authentication protocols is a composable security notion, which ensures that the scheme is secure when using it in any arbitrary environment. By using a simulator-based approach to security, several additional desirable properties to the basic functionality have been proven, such as key recycling [HLM16, Por17, GYZ17] or quantum ciphertext authentication [AGM18, DS18]. Additionally, it is possible to study the notion of authentication in the setting of computational security [BMPZ19], including public-key versions of the primitive [AGM21]. In this chapter we extensively use the abstract cryptography framework introduced by Maurer and Renner [MR11], which views cryptography as a resource theory and has been previously applied successfully to prove the composable security of purity testing based authentication schemes by Portmann [Por17].

If we define efficiency in terms of amount of qubits needed to obtain certain correctness and composable security for a constant-error quantum depolarizing channel, we show how the resulting scheme is more efficient than applying the codes separately, we include a comparison of of the threshold scheme with the Clifford and trap codes in table 4.1.

4.2 Preliminaries

Although the general quantum-information notation can be found in Chapter 2, here we will introduce some concepts and prior results that we need only in this chapter. In particular, we include a short introduction to quantum error correction and quantum message authentication.

Code	Performance	Fix $\varepsilon = \delta^{\kappa}$.
Threshold	$\max\Big\{\Omega\Big(\log(1/\varepsilon)^{1/\alpha}\Big),$ $\Omega\Big(\log(1/\delta)^{1/\alpha}\Big)\Big\}$ Thm. 4.4.3	$\Omega\Big(\log(1/\delta)^{1/\alpha}\Big)$
Trap	$\Omega\Big([\log(1/\varepsilon)\log(1/\delta)]^{1/\alpha}\Big)$ Prop. 4.3.2	$\Omega\Big(\log(1/\delta)^{2/\alpha}\Big)$
Clifford	$\Omega\Big([\log(1/\varepsilon)\log(1/\delta)]^{1/\alpha}\Big)$ Prop. 4.3.3	$\Omega\Big(\log(1/\delta)^{1/\alpha+1}\Big)$

Table 4.1: Total number of qubits for authenticated noisy channels with inner and outer $[[n, 1, d]]$-QECC of decay α.

4.2.1 Notation

We define the *Pauli group* as the group of n-fold tensor products of single-qubit Pauli matrices, that is

$$\mathcal{G}_n := \{i^k P_1 \otimes P_2 \otimes \cdots \otimes P_n \colon \text{where } P_j \in \{I, X, Z, Y\}, k \in [4]\}\,, \tag{4.2}$$

We call the elements of the Pauli group *Pauli operations*. The *weight* $\Omega(P)$ of an n-qubit Pauli operation is the number of non-identity Pauli matrices in the n-fold tensor product. Moreover, we denote by $\Omega_X(P)$ and $\Omega_Z(P)$ the number of X and Z-Pauli matrices respectively.

The *Clifford group* is the group of n-qubit unitaries that leave the Pauli group invariant. That is,

$$\mathcal{C}_n := \{C \in \mathrm{U}(2^n) \colon i^k C P C^{\dagger} \in \mathcal{G}_n \text{ for all } P \in \mathcal{G}_n, k \in [4]\}\,. \tag{4.3}$$

4.2.2 Useful lemmas

The following variant of Chernoff's bound is the probability of the majority in a population becoming the minority, and vice versa.

Lemma 4.2.1 ([GH01, Lemma 1])**.** *Given a set A, let $B \subseteq A$. Let $k \in \mathbb{N}$ be a positive integer such that $k < |A|$. Let S be a randomly sampled subset of A of size k, then for any $0 < \gamma \leq 1$ we have*

$$P\left[\frac{|S \cap B|}{k} < (1-\gamma)\frac{|B|}{|A|}\right] < \exp\left(-\gamma^2 \frac{|B|}{|A|}\frac{k}{2}\right).$$

4.2.3 Quantum error correction

Since quantum information is very sensitive to errors and noise from the environment; quantum error correction is developed as a tool to protect data against errors. A $[[n, k, d]]$ quantum error-correcting code (QECC) is an encoding of k 'logical qubits' (which we wish to protect from errors) into a codeword consisting of n 'physical qubits' (auxiliary qubits), with $n > k$. The distance d is the minimum weight of a Pauli P to convert one valid codeword into another.

After the encoded information is subjected to noise, we perform a collective measurement on the n qubits which will enable us to diagnose the type of error that occurred, called error syndrome. Afterwards, error decoding or recovery is performed, to return to the original state of the code. We say that a $[[n, k, d]]$ QECC can correct t errors if recovery is successful for any superoperator with support on the set of Pauli operators of weight up to t. In any case, we assume that we can always decode, possibly to a different state than the input if more than t errors are present. Moreover, sometimes we are satisfied just with knowing if an error has occurred, without the need to reverse it. We call this the error-detection property of the code. In fact, a QECC with distance d can correct $t = (d-1)/2$ errors. For a more in-depth analysis we refer the reader to standard literature in error correction [NC10, Pre99].

Stabilizer codes introduced by Gottesman [Got96] allow us to describe quantum states in terms of the operators stabilizing them instead of working with the state itself, by means of group theory techniques for the Pauli group. Any two elements of the Pauli group $\mathcal{G}_n$ either commute or anti-commute and square to $\pm I$, which we will use to describe codewords. Given an abelian subgroup S of $\mathcal{G}_n$, we define the *stabilizer code* V_S to be the stable states under the action of elements of S. That is,

$$V_S := \{|\psi\rangle : M|\psi\rangle = |\psi\rangle, \quad M \in S\} . \tag{4.4}$$

Let us denote by $S_1, \ldots, S_m$ the generators of the stabilizer group $S = \langle S_1, \ldots, S_m \rangle$. Since any Pauli error $P \in \mathcal{G}_n$ either commutes or anti-commutes with each element of the generator, we can define the vector of corrections $s_P = (s_{1,P}, \ldots, s_{m,P})$ such that $s_{j,P} = 0$ if P_j commutes with S_j, and $s_{j,P} = 1$ if it anti-commutes. Therefore,

$$S_j P|\psi\rangle = (-1)^{s_{j,P}} P S_i |\psi\rangle = (-1)^{s_{j,P}} P|\psi\rangle, \quad \text{for all } |\psi\rangle \in V_S . \tag{4.5}$$

We call the vector s the *syndrome* of the error-correcting code. Errors with non-zero syndrome for some element in the stabilizer $M \in S$ can be detected by the QECC – i.e. the ones that anti-commute with some element of the stabilizer. However, commuting errors are undetectable, and will change the code whenever they are not part of the stabilizer. If we denote by $S^\perp$ the set of Paulis that commute with the stabilizer, i.e.

$$S^\perp := \{P \in \mathcal{G}_n \colon PM = MP \text{ for all } M \in S\}\,, \tag{4.6}$$

then the set of undetectable errors that change the data non-trivially is $S^\perp \backslash S$.

Purity testing codes are exactly the stabilizer codes that detect any non-trivial Pauli attack with high probability. This property makes them extremely well suited for constructing authentication schemes as we will see in Section 4.2.4.

Definition 4.2.2. *A set $\{V_k\}_{k\in\mathcal{K}}$ of stabilizer codes, each with respective stabilizer subgroup S_k, is ε-*purity testing *if, when the code is selected uniformly at random, the probability of any Pauli error $P \in \mathcal{G}_n$ acting non-trivially on the data and not being detected is upper bounded by ε. That is,*

$$\Pr_{k\in\mathcal{K}}\Big(P \in S_k^\perp \setminus S_k\Big) \le \varepsilon\,.$$

There is also a more general family of codes that will be useful for our analysis.

Definition 4.2.3. *Let $\alpha \in (0,1]$. We say that a family of quantum error-correcting codes $[[n,1,d]]$ with stabilizer group S and threshold p_{th} has* decay of order *α if there exist constants $\kappa, \beta > 0$ such that*

$$\Pr\Big(\boldsymbol{X} \in S^\perp \setminus S\Big) \le \kappa (p/p_{th})^{\beta n^\alpha}\,, \quad \textit{when } p < p_{th}\,.$$

Note that the distance of a code is uniquely determined by the size $d = \Theta(n^\alpha)$.

For example, well known *concatenated* codes fall into this category. Given a $[[n,1,d]]$ QECC, we can recursively encode each encoded qubit in n physical qubits, which can be encoded again such that each layer L of concatenation is a $[[n^L, 1, d^L]]$ QECC, see [Pre99]. After L levels of concatenation, the probability of failed recovery is upper bounded by

$$\Pr\Big(\mathbf{X} \in S^\perp \setminus S\Big) \le \binom{n}{t+1}^{-1}\left(\binom{n}{t+1}p\right)^{(t+1)^L}\,. \tag{4.7}$$

Note that if $p < p_{\text{th}} := \binom{n}{t+1}^{-1}$, then we can make the failure probability as small as desired by increasing the number of layers. This is, a $[[n^L, 1, d^L]]$ concatenated codes has decay $\alpha := \log_n(t+1)$.

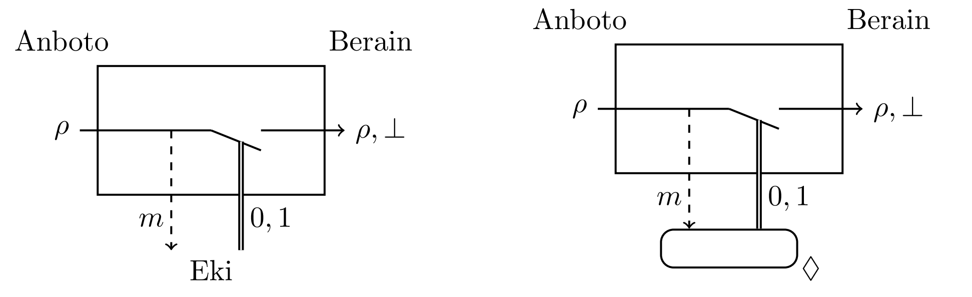

(a) Secure authenticated quantum channel with adversary Eki.

(b) Secure authenticated quantum channel with no adversary present.

Figure 4.1: Characterization of an authenticated quantum channel $\mathcal{S}_\Diamond$.

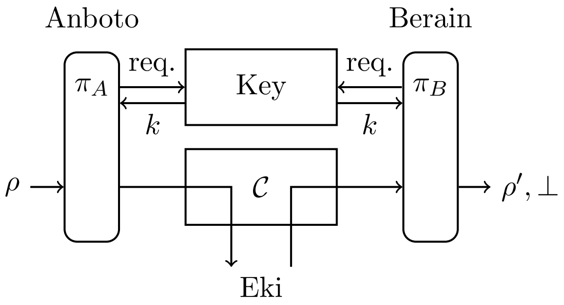

Figure 4.2: The real system for quantum message authentication.

4.2.4 Quantum authentication

In the context of constructive cryptography, a quantum authentication protocol is expected to construct an *authenticated quantum channel*, $\mathcal{S}$, from nothing but an insecure quantum channel and a secret key source. The goal of a secure quantum channel is to allow Anboto to send m qubits to Berain without Eki tampering with the data. On the one hand, they cannot stop Eki from learning that a message has been transmitted nor cutting the communication lines. Hence Eki's actions can be described as a bit 0 when Berain gets the message, and 1 when he does not. On the other hand, in the presence of no adversary, Eki's interface is substituted by a filter $\Diamond_E$ that models an honest behavior, in this case always allowing Berain to receive exactly the message that Anboto sent. A graphical description of the channel $\mathcal{S}_\Diamond$ is given in figure 4.1.

In order to construct the filtered resource $\mathcal{S}_{\Diamond}$, quantum authentication protocols will use a shared secret key $\mathcal{K}$ and an insecure quantum channel $\mathcal{C}_{\#}$, here the filter $\Diamond_E$ represents an honest behavior instead of an adversary, and the filter $\#_E$ is a noisy channel. After receiving a message ρ, the protocol π_A authenticates it with the key k received from $\mathcal{K}$ and sends the message to the insecure quantum channel $\mathcal{C}$. The protocol π_B upon receiving a message checks its validity with the shared secret key k, and outputs either ρ' or an error message $\perp$. In absence of an adversary, we substitute Eki's interface by a noise filter $\#_E$. Note that for our purposes we are not considering key resources, which greatly simplifies Portmann's descriptions [Por17, Section 3].

A generic way of constructing authentication codes was given by Barnum et al. [BCG+02] using purity-testing codes. In these schemes, the message is first encoded using a $[[n, 1, d]]$ purity-testing error-correcting code, and then encrypted with a quantum one-time pad using the shared secret key, see Protocol 5.

Protocol 5 Quantum 'encode-then-encrypt' message authentication scheme from purity-testing codes.

Setting. At the beginning of the protocol the encoder and decoder have access to secret shared key $(k, l) \in \mathcal{K}$, with $\mathcal{K} := \mathcal{K}_0 \times \mathcal{K}_1$ not necessarily of the same size. Let $\{V_k\}_{k \in \mathcal{K}_0}$ be a family of purity-testing codes.

Encoding.

1: Given input data qubit ρ_d, append the $(n-1)$-qubit syndrome state $|0\rangle\langle 0|^{\otimes(n-1)}$.

2: Encode everything with the purity testing code according to the secret shared key $k \in \mathcal{K}_0$, to obtain $V_k(\rho_d \otimes |0\rangle\langle 0|^{\otimes(n-1)})V_k^\dagger$.

3: Encrypt the message with a quantum one-time pad using the key $l \in \mathcal{K}_1$, the final message is the following

$$\mathrm{Auth}_{k,l}(\rho_d) = P_l V_k(\rho_d \otimes |0\rangle\langle 0|^{\otimes(n-1)})V_k^\dagger P_l\,. \tag{4.8}$$

Decoding.

1: First decrypt the data using l.

2: Decode the data according to the error-correcting code V_k given by the key k.

3: Measure the syndrome register in the computational basis. If the measurement outcome is 0, accept the protocol. Else, abort.

We previously defined a quantum message authentication system as a protocol that constructs an authenticated quantum channel $\mathcal{S}_{\Diamond}$ as in figure 4.1, from some shared secret key $\mathcal{K}$ and an insecure quantum channel $\mathcal{C}_{\#}$, where the filter introduces noise. Portmann showed that the scheme from Protocol 5 based on purity-testing codes provides quantum authentication protocols, given that the filter is noiseless, denoted $\square_E$.

Theorem 4.2.4 ([Por17, Lemma D.1])**.** *Given a δ-purity testing protocol $[[n,1,d]]$, let $\pi_{AB}^{auth} = (\pi_A, \pi_B)$ denote the converter corresponding to Anboto and Berain's protocols 5. Then π_{AB}^{auth} constructs an authenticated quantum channel $\mathcal{S}_{\Diamond}$, given an insecure noiseless quantum channel $\mathcal{C}_{\square}$ and a secret shared key $\mathcal{K}$ within $(0, \delta^{auth})$, where $\delta^{auth} = \max\{\delta, 2^{-(n-1)}\}$. That is,*

$$\mathcal{C}_{\square} || \mathcal{K} \xrightarrow{\pi_{AB}^{auth}, (0, \delta^{auth})} \mathcal{S}_{\Diamond}\,.$$

4.2.5 Noisy channels

Given a noisy quantum channel between Anboto A and Berain B, where the *noise* is represented by a quantum operation $\mathcal{F}_{A\to B}$, we say that there exists an error-correction protocol π_{AB}^{ecc}, defined by an encoding map $\mathcal{E}_A$ and a decoding map $\mathcal{D}_B$, correcting the errors induced by $\mathcal{F}_{A\to B}$ within ε^{ecc}, if

$$\frac{1}{2} \|\mathcal{D}_B \circ \mathcal{F}_{A\to B} \circ \mathcal{E}_A - I_{A\to B}\|_{\diamond} \leq \varepsilon^{\text{ecc}}. \tag{4.9}$$

We can rewrite the above statement in the abstract cryptography language.

Lemma 4.2.5 ([Por17, Lemma 4.2])**.** *Let $\#_E$ be a filter introducing the noise given by the quantum operation $\mathcal{F}$, and let $\square_E$ be a noiseless filter. If there exists an error correction protocol $\pi_{AB}^{\text{ecc}} = (\pi_A^{\text{ecc}}, \pi_B^{\text{ecc}})$ that corrects the errors induced by $\mathcal{F}$ within ε^{ecc}, then π_{AB}^{ecc} constructs a noiseless channel $\mathcal{C}_{\square}$, from a noisy channel $\hat{\mathcal{C}}_{\#}$ within $(\varepsilon^{\text{ecc}}, 0)$. That is,*

$$\hat{\mathcal{C}}_{\#} \xrightarrow{\pi_{AB}^{\text{ecc}}, (\varepsilon^{ecc}, 0)} \mathcal{C}_{\square}\,.$$

It is now clear what the relevance of splitting the cryptographic security definition in terms of correctness and security is, a direct consequence of Theorems 2.2.5 and 4.2.4, and Lemma 4.2.5, is that a δ-secure authentication scheme wrapped in an ε-correct error-correcting code constructs an (ε, δ)-secure authenticated quantum channel, instead of an $(\varepsilon+\delta)$-secure from the original composition theorem.

In order to be able to compare explicit schemes, we will restrict to a basic type of noise, typically used in error correction literature, i.i.d. Pauli noise. We will assume that when qubits are sent through a noisy channel, they independently undergo a X, Y or Z Pauli error with probabilities p_X, p_Y and p_Z respectively. This model is interesting not only because it models many interesting real situations, but also because the Pauli operators are a basis of single-qubit operations, and thus protection against i.i.d. Pauli noise for a single qubit implies protection against any single-qubit error. Moreover, in the 'encode-then-encrypt' authentication scheme, the one-time-pad encryption and the Pauli twirl make any attack become a Pauli attack. This means that for the security proof it is enough to prove security against Pauli attacks [BGS13], which is also the reason why the underlying error-correction codes are required to be purity-testing codes.

Let $\mathcal{F}_n$ denote a quantum noise channel acting on n-qubits independently, which we can write in terms of the basis elements of single-qubit operations $\mathcal{F}_n = \mathcal{F}^{\otimes n}$, where

$$\mathcal{F}(\rho) := (1 - p_X - p_Y - p_Z)\rho + p_X X\rho X^\dagger + p_Y Y\rho Y^\dagger + p_Z Z\rho Z^\dagger \,. \qquad (4.10)$$

This channel leaves the state untouched with probability $p_I := 1 - p_X - p_Y - p_Z$ and each Pauli operation is applied with probability p_X, p_Y and p_Z respectively. As mentioned earlier, we will assume that the noise is i.i.d. distributed. Therefore, we can write the noise channel acting on a n-qubit register as

$$\mathcal{F}_n(\rho) = \sum_{\substack{k_1+k_2+k_3+k_4=n \\ k_1,k_2,k_3,k_4\geq 0}} \binom{n}{k_1, k_2, k_3, k_4} \prod_{Q\in\{I,X,Y,Z\}} p_j^{k_j} Q^{k_Q} \rho (Q_Q^\dagger)^{k_Q} \,. \qquad (4.11)$$

The depolarizing channel, the most commonly used noise model in error-correction literature [Ter15], is of this type. When a qubit goes through the depolarizing channel, the channel erases the qubit and substitutes it by a completely mixed state $I/2$ with probability p, and leaves the qubit untouched with probability $1 - p$. In the notation from eq. (4.11), this is the same as saying that with probability $1 - p$ the qubit is being left untouched, and each Pauli operation will be applied with probability $p/3$. Therefore, the depolarizing noise acting on n-qubits can be written as

$$\begin{aligned}\mathcal{F}_n(\rho) &= \sum_{\substack{k_1+\cdots+k_4=n\\ k_1,\ldots,k_4\geq 0}} \binom{n}{k_1,\ldots,k_4}(1-p)^{k_1}\left(\frac{p}{3}\right)^{n-k_1} \prod_{Q\in\{X,Y,Z\}} Q^{k_Q}\rho(Q_Q^{\dagger})^{k_Q},\\ &= (1-p)^n\rho + (1-p)^{n-1}\frac{p}{3}\sum_{k=1}^{n}\sum_{Q\in\{I,X,Y,Z\}} Q^k\rho(Q^{\dagger})^k + \cdots\,. \end{aligned} \tag{4.12}$$

4.3 Explicit composed protocols

Attempting to protect our data against both noise and attacks can be seen in the abstract cryptography (AC) framework of Section 2.2.2 as constructing a noiseless secure channel $\mathcal{C}$ (correct and secure) from a noisy insecure channel $\hat{\mathcal{C}}_{\#}$. In practice, this is obtained by concatenating an authenticating scheme from Section 4.2.4 with an error-correcting code from Section 4.2.3. Unfortunately, both these constructions involve encoding logical qubits in redundant physical qubits for protection, and thus the dimension of the secure and insecure channels can differ vastly.

In this section we analyse the cost-effectiveness of two of the most used authentication schemes: the trap and Clifford schemes. They are both purity testing code families, see Definition 4.2.2, i.e. the probability of non-trivial errors being undetected by the codes is low. In short, the trap scheme consists of two error-correcting codes, such that the inner code corrects low-weight noise and the traps detect high-weight attacks, hence the adversary cannot effectively choose a relevant attack without being detected. The Clifford scheme on the other hand is a single error-correcting code that randomizes the weight of the error, actually making it *strong* purity testing, meaning it detects *any* error with high probability.

Taking the amount of qubits necessary as a parameter for efficiency, in this section we analyse the cost-effectiveness of protecting a single qubit of data against both noise and attacks. We study both the trap of the Clifford scheme for the authentication and consider only single-qubit error-correcting codes.

4.3.1 Trap scheme

Given a fixed $[[n,1,d]]$ error-correcting code, defined by encoder $\mathsf{Enc}_n\colon \mathbb{C}^2 \to \mathbb{C}^{2^n}$ and decoder $\mathsf{Dec}_n\colon \mathbb{C}^{2^n} \to \mathbb{C}^2$, the trap authentication scheme constructs a set $\{V_k\}_{k\in\mathcal{K}}$ of purity testing codes encoding 1 logical qubit in $3n$ physical qubits, by first encoding each data qubit ρ_d in n physical qubits, appending $2n$ 'traps' to the encoded data (n copies of $|0\rangle\langle 0|$ and another n copies of $|+\rangle\langle +|$); the resulting $3n$-qubit registers are permuted attending to a secret shared key $k \in \mathcal{K}$. The purity-testing codes needed for the encode-then-encrypt scheme in Protocol 5 are given by:

$$\begin{aligned} V_k(\rho_d) &= \pi_k(\mathsf{Enc}_n(\rho_d) \otimes |0\rangle\langle 0|^n \otimes |+\rangle\langle +|^n)\pi_k^\dagger \\ &:= \pi_k(I_{2n} \otimes H^{\otimes n})(\mathsf{Enc}_n(\rho_d) \otimes |0\rangle\langle 0|^{2n-1})(I_{2n} \otimes H^{\otimes n})\pi_k^\dagger \,. \end{aligned} \tag{4.13}$$

When decoding, first the inverse permutation, according to the secret key, is applied. Finally, the data registers are decoded according to the fixed error-correcting code and the traps are measured in the computational and Hadamard bases respectively. Here we consider the underlying code as error *correcting*, for the sake of fair comparison with later codes, but there is also an error *detection* variant of the trap code [BGS13].

Lemma 4.3.1 ([BW16, Theorem 5.2])**.** *The trap code from eq. (4.13) with inner error-correcting code* $[[n, 1, d]]$ *is* $(1/3)^{\frac{d+1}{2}}$*-purity testing.*

Given a purity testing code, the discussion in Sections 4.2.4 and 4.2.5 ensures us that if there exists an error-correcting protocol correcting the errors induced by the noisy channel within error ε^{ecc}, then composing qubit-wise the trap code encoded qubits with the error-correcting code will give rise to an ε^{ecc}-correct and $(1/3)^{\frac{d+1}{2}}$-secure protocol. The following theorem rephrases the security and correctness of the composed protocol in terms of the number of qubits. For simplicity of the analysis we assume the channel is affected by depolarizing noise.

Proposition 4.3.2. *Let* π^{trap} *be the trap authentication scheme with a family* $[[n_{\text{in}}, 1, d_{\text{in}}]]$ *of inner codes of decay* α_{in}*, and* π^{ecc} *a family* $[[n_{\text{out}}, 1, d_{\text{out}}]]$ *of outer codes of decay* α_{out}*, with* $\mathfrak{p}_{\text{in}}$ *and* $\mathfrak{p}_{\text{out}}$ *thresholds respectively and* $\kappa_{\text{in}} < \mathfrak{p}_{\text{out}}$*. Let* $\#_E$ *be a filter introducing the noise given by the depolarizing channel with channel error* $p < \mathfrak{p}_{\text{in}}$ *and* $p < \mathfrak{p}_{\text{out}}$*. Then to obtain* ε*-correctness and* δ*-security, i.e.*

$$\hat{\mathcal{C}}_{\#}||\mathcal{K} \xrightarrow{\pi^{trap}\pi^{ecc},(\varepsilon,\delta)} \mathcal{S}_{\Diamond}^m \,, \tag{4.14}$$

it is sufficient for the total amount of qubits n_{total} *to grow as*

$$n_{total} = \Omega\Big(\log(1/\varepsilon)^{1/\alpha_{\text{out}}} \log(1/\delta)^{1/\alpha_{\text{in}}}\Big) \,. \tag{4.15}$$

Proof. Let us denote by S_{in} the stabilizer subgroup of the inner code and by S the one of the trap code concatenated with the error-correcting code.

Security. Since the inner code uniquely determines the security it is natural to start with it. By Lemma 4.3.1, given $P \in \mathcal{G}_{n_{\text{in}}}$, δ-security is obtained whenever

$$\Pr_{k\in\mathcal{K}}\Big(P \in (S_{\text{in}})_k^{\perp} \setminus (S_{\text{in}})_k\Big) \le \left(\frac{1}{3}\right)^{\frac{\Theta\left(n_{\text{in}}^{\alpha_{\text{in}}}\right)+1}{2}} \le \delta\,. \tag{4.16}$$

Taking logarithms on both sides we obtain,

$$\frac{\Theta\big(n_{\text{in}}^{\alpha_{\text{in}}}\big)+1}{2} \ge \frac{\log(1/\delta)}{\log(3)} \quad \text{i.e.} \quad n_{\text{in}} \ge \Omega\Big(\log(1/\delta)^{1/\alpha_{\text{in}}}\Big)\,. \tag{4.17}$$

Correctness. Note in eq. (4.13) how the first $n_{\text{in}}n_{\text{out}}$ qubits have double encoding, and the last $2n_{\text{in}}n_{\text{out}}$ only one, thus the probability of being detected for $n_{\text{in}}n_{\text{out}}$ of the qubits is not uniquely determined by the outer error-correcting code. Then the error probability is bounded by

$$\begin{aligned}\Pr\Big(\mathbf{X} \in S^{\perp} \setminus S\Big) &\le \kappa_{\text{out}}\Big(\Pr\Big(\mathbf{X} \in S_{\text{in}}^{\perp} \setminus S_{\text{in}}\Big)/p_{\text{out}}\Big)^{n^{\alpha_{\text{out}}}} \\ &\quad + 2n_{in}\kappa_{\text{out}}(p/p_{\text{out}})^{n^{\alpha_{\text{out}}}} \qquad (4.18)\\ &\le \kappa_{\text{out}}\left(\frac{\kappa_{\text{in}}}{p_{\text{out}}}\right)^{n_{\text{out}}^{\alpha_{\text{out}}}}\left(\frac{p}{p_{\text{in}}}\right)^{n_{\text{in}}^{\alpha_{\text{in}}}n_{\text{out}}^{\alpha_{\text{out}}}} + 2n_{in}\kappa_{\text{out}}\left(\frac{p}{p_{\text{out}}}\right)^{n_{\text{out}}^{\alpha_{\text{out}}}}, \qquad (4.19)\end{aligned}$$

thus we achieve ε-correctness whenever $\Pr\big(\mathbf{X} \in S^{\perp} \setminus S\big) \le \varepsilon$, thus it would be enough for both following conditions to hold

$$\begin{cases}\kappa_{\text{out}}\left(\frac{\kappa_{\text{in}}}{p_{\text{out}}}\right)^{n_{\text{out}}^{\alpha_{\text{out}}}}\left(\frac{p}{p_{\text{in}}}\right)^{n_{\text{in}}^{\alpha_{\text{in}}}n_{\text{out}}^{\alpha_{\text{out}}}} \le \varepsilon/2 \\ 2n_{in}\kappa_{\text{out}}\left(\frac{p}{p_{\text{out}}}\right)^{n_{\text{out}}^{\alpha_{\text{out}}}} \le \varepsilon/2\,.\end{cases} \tag{4.20}$$

For the first item in eq. (4.20), by taking logarithms we need

$$n_{\text{out}}^{\alpha_{\text{out}}}\log\left(\frac{p_{\text{out}}}{\kappa_{\text{in}}}\right) + n_{\text{in}}^{\alpha_{\text{in}}}n_{\text{out}}^{\alpha_{\text{out}}}\log\frac{p_{\text{in}}}{p} \ge \log\left(\frac{2\kappa_{\text{out}}}{\varepsilon}\right), \tag{4.21}$$

thus it is enough for any of the following two to hold,

$$n_{\text{out}}^{\alpha_{\text{out}}}\log\left(\frac{p_{\text{out}}}{\kappa_{\text{in}}}\right) \ge \log\left(\frac{2\kappa_{\text{out}}}{\varepsilon}\right) \quad \text{and} \quad n_{\text{in}}^{\alpha_{\text{in}}}n_{\text{out}}^{\alpha_{\text{out}}}\log\frac{p_{\text{in}}}{p} \ge \log\left(\frac{2\kappa_{\text{out}}}{\varepsilon}\right). \tag{4.22}$$

Note that the first requirement is weaker than the second one, thus it is enough to ask

$$n_{\text{out}} \geq \Omega\Big(\log(1/\varepsilon)^{1/\alpha_{\text{out}}}\Big)\,. \tag{4.23}$$

For the second item in eq. (4.20), by taking logarithms we need

$$n_{\text{out}}^{\alpha_{\text{out}}} \log\left(\frac{p_{\text{out}}}{p}\right) \geq \log\left(\frac{4n_{\text{in}}\kappa_{\text{out}}}{\varepsilon}\right) \tag{4.24}$$

that is,

$$n_{\text{out}} \geq \Omega\left(\left[\log(n_{\text{in}}) + \log\left(\frac{1}{\varepsilon}\right)\right]^{1/\alpha_{\text{out}}}\right)\,, \tag{4.25}$$

but this requirement is weaker than the one in eq. (4.23), thus proving the desired bound. □

4.3.2 Clifford scheme

The Clifford authentication scheme is based on a set of purity testing unitaries $\{C_k\}_{k\in\mathcal{K}}$ given by the Clifford group. The authentication scheme first appends to the data qubit ρ_d, n 'traps' (n copies of $|0\rangle\langle 0|$) and finally a Clifford element C_k is applied to the resulting $(n+1)$-qubit registers attending to a secret shared key $k \in \mathcal{K}$. Without the need for encryption, the authenticated data is directly given by:

$$\text{Auth}_k(\rho_d) := C_k(\rho_d \otimes |0\rangle\langle 0|^n)C_k^\dagger\,. \tag{4.26}$$

When decoding, first the inverse Clifford operation, according to the secret key, is applied. Finally, the traps are measured in the computational basis.

It is not difficult to see that the n-trap Clifford authentication scheme is 2^{-n}-secure. However, the Clifford twirl will map any Pauli operation to an arbitrary-weight one, not being able to distinguish between low and high weight operations and hence making it impractical over noisy channels. It is therefore imperative to compose it with an error-correcting code for practical uses.

Proposition 4.3.3. *Let π^c be the n_{in}-qubit Clifford authentication scheme, and π^{ecc} a family $[[n_{\text{out}}, 1, d_{\text{out}}]]$ of outer codes of decay α_{out}, with p_{out} threshold. Let $\#_E$ be a filter introducing the noise given by the depolarizing channel with channel error $p < p_{\text{out}}$. Then to obtain ε-correctness and δ-security, i.e.*

$$\hat{\mathcal{C}}_{\#}||\mathcal{K} \xrightarrow{\pi^c\pi^{ecc},(\varepsilon,\delta)} \mathcal{S}_{\Diamond}^m\,, \tag{4.27}$$

it is sufficient for the total amount of qubits n_{total} to grow as

$$n_{total} = \Omega\Big(\log(1/\varepsilon)^{1/\alpha_{\text{out}}} \log(1/\delta)\Big)\,. \tag{4.28}$$

Proof. Let us denote by S_c the stabilizer subgroup of the Clifford code and by S the one of the Clifford code concatenated with the error-correcting code.

Security. Since the n_{in}-trap Clifford code is $2^{-n_{\text{in}}}$ purity-testing, we obtain δ-security whenever

$$\Pr_{k\in\mathcal{K}}\Big(P \in (S_c)_k^{\perp} \setminus (S_c)_k\Big) \leq 2^{-n_{\text{in}}} \leq \delta \quad \text{i.e.} \quad n_{\text{in}} \geq \Omega(\log(1/\delta))\,. \tag{4.29}$$

Correctness. Note that none of the errors can permeate the outer error-correcting code, because the Clifford operation will map it to an arbitrary weight Pauli operation and therefore will be detected by the traps. Hence we achieve ε-correctness whenever

$$\Pr\Big(\mathbf{X} \in S^{\perp} \setminus S\Big) \leq (n_{\text{in}}+1)\kappa_{\text{out}}\left(\frac{p}{p_{\text{out}}}\right)^{n_{\text{out}}^{\alpha_{\text{out}}}} \leq \varepsilon\,. \tag{4.30}$$

Taking logarithms on both sides we obtain,

$$n_{\text{out}}^{\alpha_{\text{out}}} \geq \log\left(\frac{p_{\text{out}}}{p}\right)^{-1} \log\left(\frac{(n_{\text{in}}+1)\kappa_{\text{out}}}{\varepsilon}\right) \quad \text{i.e.} \quad n_{\text{out}} \geq \Omega\Big(\log(1/\varepsilon)^{1/\alpha_{\text{out}}}\Big)\,. \tag{4.31}$$

□

4.4 The threshold authentication scheme

Both Hayden, Leung and Mayers [HLM16] and Portmann [Por17] constructions of composed protocols, it is assumed that the authentication scheme rejects whenever an error is present – which is always the case with very high probability when sending information through noisy channels – and therefore an error-correcting code is necessary to make the schemes useful. However, from the structure of the composition, the number of qubits used in such a construction blows up both with the size of the purity-testing code used in the authentication scheme and the error-correcting code. It is therefore natural to ask if such a composition is even necessary, and if we cannot design a protocol that directly constructs an authenticated quantum channel from a noisy insecure channel and shared secret key. This is exactly what the threshold scheme we propose in this section does.

The threshold scheme is an adaptation of the trap scheme where, with the same encoding, we require Berain to accept the message whenever *low* amount of errors are detected. In other words, we use the traps as they were originally intended, to measure the amount of error present in the encoded data and decide if these errors pertain to noise or an attack. In principle, this should not be enough, as the correctable errors of an error-correcting code grow sub-linearly in the size of the code, while for example for the depolarizing channel the number of errors is linear in the size of the code. However, error-correcting codes with fixed decay actually correct linear amount of errors with *very high probability*, which is enough to form a purity-testing family of codes.

The threshold code is constructed as follows. Given a fixed $[[n,1,d]]$ error-correcting code with encoder $\mathsf{Enc}_n\colon \mathbb{C}^2 \to \mathbb{C}^{2^n}$ and decoder $\mathsf{Dec}_n\colon \mathbb{C}^{2^n} \to \mathbb{C}^2$, the threshold authentication scheme first encodes each data qubit ρ_d in n physical qubits and then appending $2n$ 'traps' to the encoded data (n copies of $|0\rangle\langle 0|$ and another n copies of $|+\rangle\langle +|$); the resulting $3n$-qubit registers are permuted attending to a secret shared key $k \in \mathcal{K}$. The authenticated states are given by:

$$V_k(\rho_d) := \pi_k(\mathsf{Enc}_n(\rho_d) \otimes |0\rangle\langle 0|^n \otimes |+\rangle\langle +|^n)\pi_k^\dagger\,. \tag{4.32}$$

When decoding, first the inverse permutation, according to the secret key, is applied. Finally, the data registers are decoded according to the fixed error-correcting code and the traps are measured in the computational and Hadamard bases respectively. However, *only if less than a threshold* τn, with $\tau \in [0,2]$, of errors are present does the receiver accept the message. The threshold τn allows us to tune the the amount of error we are willing to accept, depending on the noise of the channel, more efficiently than adding an entire new error-correcting code to each qubit. The explicit construction of the threshold authentication scheme is given in Protocol 6.

Protocol 6 Threshold authentication scheme π_{AB}^{thr}.

Setting. At the beginning of the protocol the encoder and decoder have access to secret shared key $(k,l) \in \mathcal{K}$, with $\mathcal{K} := \mathcal{K}_0 \times \mathcal{K}_1$ not necessarily of the same size. Let $(\mathsf{Enc}_n, \mathsf{Dec}_n)$ be an $[[n,1,d]]$ error-correcting code and $\{\pi_k\}_{k\in\mathcal{K}_0}$ a family of permutations.

Encoding.

```
1: Encode the input data qubit into $n$ qubits with the fixed error
   correcting code $\mathsf{Enc}_n(\rho_d)$.
2: Append $2n$ computational basis states $|0\rangle\langle 0|^{\otimes 2n}$ and apply a Hadamard
   gate to each of the last $n$ qubits.
3: Apply the permutation $\pi_k$ to all the qubit registers according to the
   secret key $k \in \mathcal{K}_0$.
4: Finally, encrypt the message with a quantum one-time pad using the
   secret key $l \in \mathcal{K}_1$, obtaining thus
```

$$\mathrm{Auth}_{k,l}(\rho_d) = P_l \pi_k \big(\mathrm{Enc}_n(\rho_d) \otimes |0\rangle\langle 0|^{\otimes n} \otimes |+\rangle\langle +|^{\otimes n}\big) \pi_k^\dagger P_l \,. \tag{4.33}$$

Decoding.

```
1: First decrypt the data using $l$.
2: Apply the inverse permutation according to $k$.
3: Measure the second to last $n$ registers in the computational basis, and
   the last $n$ registers in the Hadamard basis. If less than a threshold $\tau n$
   of qubits differ from the expected outcome, i.e. $|0\rangle\langle 0|^n \otimes |+\rangle\langle +|^n$, accept
   the protocol. Else, abort.
4: If the protocol has been accepted, decode the data register with the
   decoder $\mathsf{Dec}_n$.
```

Since there is no outer error-correcting code in our protocol, we have to ensure that the threshold scheme constructs directly a noiseless secure quantum channel from nothing but a noisy insecure quantum channel and a shared secret key. We will separate this task in two steps, first proving the correctness and then the security.

4.4.1 Correctness

We want to prove that when we use our protocol with a noisy channel, the outcome is nearly indistinguishable from using a noiseless secure channel without adversary, see figure 4.3.

Proposition 4.4.1. *Let $\#_E$ be a filter introducing the noise given by the depolarizing channel with channel error p. Let $[[n,1,d]]$ be a family of error-correcting codes with threshold $p_{th} > p$ and decay α. Let π_{AB}^{thr} be the threshold authentication scheme built from these error-correcting codes and with threshold parameter $\tau > 4p/3$ as in Protocol 6. Then π_{AB}^{thr} is ε-correct, i.e.*

$$d(\pi_{AB}^{thr}(\mathcal{C}_{\#} \| \mathcal{K}), \mathcal{S}_{\Diamond}) \leq \varepsilon \,, \tag{4.34}$$

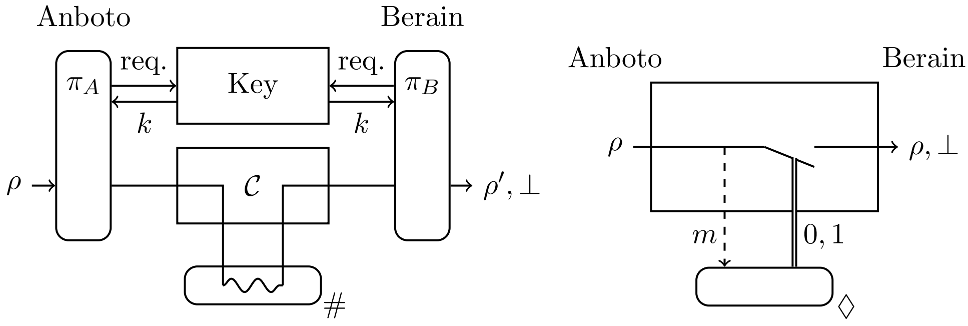

(a) Threshold protocol with no adversary present.

(b) Authenticated quantum channel with no adversary present.

Figure 4.3: Comparison between the threshold protocol in a noisy channel and a secure authenticated quantum channel without adversary.

with

$$\varepsilon = \kappa 2^{-\beta n^{\alpha} \log(p_{th}/p)} + 2^{-n(\tau - 4p/3)^2 \log(e)} \,. \tag{4.35}$$

Proof. To prove correctness within ε we have to show that the threshold protocol π_{AB}^{thr} constructs a noiseless secure channel $\mathcal{S}_{\Diamond}$ such that the real system transmitted through a noisy channel $\pi_{AB}^{\mathrm{thr}}(\mathcal{C}_{\#}||\mathcal{K})$ cannot be distinguished from the ideal system $\mathcal{S}_{\Diamond}$. Note that distinguishability in presence of no adversary is exactly the diamond norm between the identity map and the encoding-noise-decoding map of the threshold code, i.e.

$$d(\pi_{AB}^{\mathrm{thr}}(\mathcal{C}_{\#}||\mathcal{K}), \mathcal{S}_{\Diamond}) = \frac{1}{2} \left\| \mathcal{D}^{\mathrm{thr}} \circ \mathcal{F} \circ \mathcal{E}^{\mathrm{thr}} - I \right\|_{\diamond} . \tag{4.36}$$

Given a random variable $X \in \{I, X, Z, Y\}$, let $\omega(X) \in \{0, 1\}$ be a random variable denoting if the operator describes a non-trivial error on a qubit or not, i.e.

$$\omega(X_j) = \begin{cases} 1 & \text{if } X_j \neq I \\ 0 & \text{if } X_j = I \,. \end{cases} \tag{4.37}$$

Let $X_1, \ldots, X_{3n}$ be independent random variables such that the first n fail with probability p and the last $2n$ fail with probability $2p/3$, i.e.

$$\Pr(\Omega(X_j) = 1) = \begin{cases} p & \text{for } j = 1, \ldots, n, \\ 2p/3 & \text{for } j = n+1, \ldots, 3n \,. \end{cases} \tag{4.38}$$

We do this distinction because we have computational and Hadamard bases traps, thus the probability of rejection is different. We denote by $\mathbf{X}$ the tensor product of the first n variables, i.e. $\mathbf{X} := X_1, \otimes \cdots \otimes X_n$. We achieve ε-correctness whenever the rejection probability of the traps or the failed recovery of the error-correcting code encoding the data qubits is less than ε. That is,

$$\Pr\left(\left\{\mathbf{X} \in S^{\perp} \setminus S\right\} \cup \left\{\sum_{j=n}^{3n} \omega(X_j) \geq \tau n\right\}\right) \tag{4.39}$$

$$= \Pr\Big(\mathbf{X} \in S^{\perp} \setminus S\Big) + \Pr\left(\sum_{j=n}^{3n} \omega(X_j) \geq \tau n\right) \tag{4.40}$$

$$\leq \kappa(p/p_{\text{th}})^{\beta n^{\alpha}} + \exp\big(-n(\tau - 4p/3)^2\big)\,, \tag{4.41}$$

whenever $p < p_{\text{th}}$, where we used Hoeffding's inequality with $\tau > 4p/3$. □

4.4.2 Security

With security we mean that, in presence of a malicious player, there exists a simulator in the ideal protocol that is indistinguishable from the real protocol. However, instead of constructing this simulator, it is enough to show that the threshold scheme constructs a set of codes that is purity testing, which will provide us with security by Theorem 4.2.4. Although Portmann's original proof constructs a secure channel from a noiseless channel, in the security proof the filters are substituted by an adversary, and therefore work for our setting as well.

By setting the threshold properly, we can leverage the fact that error-correcting codes correct a linear amount of errors with high probability to prove that the threshold scheme is purity testing.

Proposition 4.4.2. *Let $\#_E$ be a filter introducing the noise given by the depolarizing channel with channel error p. Let $[[n, 1, d]]$ be a family of error-correcting codes with threshold $p_{th} > p$ and decay α. Let π_{AB}^{thr} be the threshold authentication scheme built from these error-correcting codes and with threshold parameter $\tau < 2p_{th}$ as in Protocol 6. Then π_{AB}^{thr} is δ-secure, i.e. there exists a simulator σ_E, δ-close to the real protocol*

$$d(\pi_{AB}^{rm}\mathcal{C}, \sigma_E\mathcal{S}) \leq \delta\,, \tag{4.42}$$

with

$$\delta = \max\left\{ \frac{10\kappa}{9\sqrt{2\pi}\sqrt{3b(1-b)}} \exp\left(-\frac{\ln(n)}{2} - \beta n^{\alpha} \ln(p_{th}/b)\right), \right.$$
$$\left. \exp\left(-n\frac{3b}{2}\left(1-\frac{\tau}{2b}\right)^2\right)\right\}, \quad (4.43)$$

where $b := \frac{p_{th}}{2} + \frac{\tau}{4}$.

Proof. We will prove security by showing that the threshold scheme constructs a family of purity-testing codes, this is, that if the permutation key is selected uniformly at random, the probability of any Pauli error $E \in \mathcal{G}_n$ acting non-trivially on the data and not being detected is upper bounded by δ. The threshold code, for a key $k \in \mathcal{K}_0$, is characterized by the codes

$$V_k(\rho_d) = \pi_k(\mathsf{Enc}_n(\rho_d) \otimes |0\rangle\langle 0|^n \otimes |+\rangle\langle +|^n)\pi_k^{\dagger}. \quad (4.44)$$

The first n qubits are used to decode the inner error-correcting code, and the last $2n$ are the traps, such that the protocol rejects whenever more than τn non-zero traps are detected. Let us denote by S_{in} the stabilizer subgroup of the inner error-correcting code. For a particular permutation π_k, on the one hand, the set of Paulis that are not detected is

$$S_k^{\perp} := \{\pi_k^{\dagger}(P \otimes Q \otimes R)\pi_k \colon P \in S_{\text{in}}^{\perp},\, \omega_X(Q) + \omega_Z(R) \leq r\}. \quad (4.45)$$

On the other hand, since the traps are invariant to Z and X operations respectively, the Paulis that act trivially on the message are

$$S_k := \{\pi_k^{\dagger}(P \otimes Q \otimes R)\pi_k \colon P \in S_{\text{in}},\, Q \in \{I, Z\}^{\otimes n},\, R \in \{I, X\}^{\otimes n}\}. \quad (4.46)$$

We can also split the set of permutations in terms of the error correction and the trap detection

$$\Pi^0(E) := \{\pi \in \Pi_{3n} \colon E = \pi^{\dagger}(P \otimes T)\pi,\, P \in S_{\text{in}}^{\perp} \setminus S_{\text{in}},\, T \in \mathcal{G}_{2n}\}, \quad (4.47)$$
$$\Pi^1(E) := \{\pi \in \Pi_{3n} \colon E = \pi^{\dagger}(P \otimes Q \otimes R)\pi,\, P \in \mathcal{G}_n,\, \omega_X(Q) + \omega_Z(R) \leq \tau n\}, \quad (4.48)$$

so that we can bound the purity testing parameter by the minimum size of both sets

$$\Pr_{k \in \mathcal{K}_0}\left(E \in S_k^{\perp} \setminus S_k\right) \leq \min\left\{\frac{|\Pi^0(E)|}{|\Pi_{3n}|}, \frac{|\Pi^1(E)|}{|\Pi_{3n}|}\right\}. \quad (4.49)$$

Given an operator $E \in \mathcal{G}_{3n}$, we will divide the proof in two cases attending to its weight, where the weight of an operator is defined as the sum of the weight of the elements in its tensor product,

$$\omega(E) := \sum_{j=1}^{3n} \omega(E_j), \quad \text{where} \quad \omega(E_j) := \begin{cases} 1 & \text{if } E_j \neq I \\ 0 & \text{if } E_j = I\,. \end{cases} \tag{4.50}$$

Note that it is enough to bound one of the sets from eq. (4.49) for different weight attacks.

Case 1: $\omega(E) \leq 3np_{\mathbf{th}}$**.** For low-weight attacks, still linear in the total size of the protocol, we expect the error-correcting code to correct them with high probability. Since the set of Pauli operators acting non-trivially and being undetected is exactly the same as the set of operators that the error-correcting code fails to decode correctly and they are randomized because of the one-time pad, we can rewrite it in terms of random variables. We define a set of i.i.d. random variables $X_1, \ldots, X_{3n}$ such that

$$\Pr(\omega(X_j) = 1) = \frac{\omega(E)}{3n} \quad \text{for } j \in [n]\,. \tag{4.51}$$

We denote by $\mathbf{X}$ the tensor product of the first n variables, that is, $\mathbf{X} := X_1, \otimes \cdots \otimes X_n$. Condition on a fixed amount of registers suffering an error, we have the bound

$$\frac{|\Pi^0(E)|}{|\Pi_{3n}|} \leq \Pr\left(\mathbf{X} \in S_{\text{in}}^{\perp} \setminus S_{\text{in}} \middle| \sum_{j=1}^{3n} \omega(X_j) = \omega(E)\right) \tag{4.52}$$

$$\leq \frac{\Pr\big(\mathbf{X} \in S_{\text{in}}^{\perp} \setminus S_{\text{in}}\big)}{\Pr\Big(\sum_{j=1}^{3n} \omega(X_j) = \omega(E)\Big)} \tag{4.53}$$

$$\leq \frac{10\kappa}{9\sqrt{2\pi}(3np_{th})^{\beta n^{\alpha}}} \frac{\omega(E)^{\beta n^{\alpha}}}{\sqrt{\omega(E)(1 - \omega(E)/3n)}}\,. \tag{4.54}$$

Case 2: $\omega(E) \geq 3\tau n/2$**.** High weight attacks will be detected by the traps with high probability, even when a linear amount of them τn are triggered before aborting the protocol. Although we cannot leverage the independence of errors as in the security proof of the trap code, we can apply a sampling variant of the Chernoff bound, see Lemma 4.2.1. Let us define the total population to be all the registers $A := \{1, \ldots, 3n\}$, and the sub-population the traps $B := \{n, \ldots, 3n\}$. Given a Pauli attack E of weight $\omega(E)$ and a random sample $S \subset A$, with $|S| = \omega(E)$, we have

$$\frac{|\Pi^1(E)|}{|\Pi_{3n}|} \leq \Pr\left(\sum_{j=n}^{3n} \omega(E_j) < \tau n\right) = \Pr\left(\sum_{j=n}^{3n} \omega(E_j) < (1-\gamma)\frac{|B|}{|A|}\omega\right) \tag{4.55}$$

$$< \exp\left(-\gamma^2 \frac{|B|}{|A|}\frac{\omega(E)}{2}\right) = \exp\left(-\frac{\omega(E)}{3}\left(1 - \frac{3\tau n}{2\Omega(E)}\right)^2\right), \tag{4.56}$$

where $\gamma = 1 - \frac{3\tau n}{2\omega}$, with $\gamma \in (0,1)$ whenever $\omega(E) > 3\tau n/2$. Finally, since the adversary will pick the optimal weight, we need a bound independent of the weight $\omega(E)$. We can ensure this by choosing a non-empty overlap of the two cases, which holds whenever $\tau < 2p_{th}$, and picking a weight in the middle of the accepted ones e.g. $\omega(E) = 3n\left(\frac{p_{th}}{2} + \frac{\tau}{4}\right) = 3nb$. This provides an upper bound on the purity-testing parameter

$$\begin{aligned}\Pr_{k \in \mathcal{K}_0}&\left(E \in S_k^\perp \setminus S_k\right) \\ &\leq \max\left\{\frac{10\kappa}{9\sqrt{2\pi}\sqrt{3b(1-b)}} \exp\left(-\frac{\ln(n)}{2} - \beta n^\alpha \ln(p_{th}/b)\right),\right. \\ &\qquad\qquad \left.\exp\left(-n\frac{3b}{2}\left(1 - \frac{\tau}{2b}\right)^2\right)\right\}. \end{aligned} \tag{4.57}$$

□

4.4.3 Efficiency in terms of qubits

We can combine the correctness and security requirements of the threshold scheme to obtain the sufficient amount of qubits that the threshold scheme requires to obtain (ε, δ)-security. The effectiveness of the threshold scheme lies in the fact that instead of encoding the traps in an error-correcting code to allow some error, we can just tune the threshold parameter τ to the noise scenario whilst being able to detect adversaries. That is, in contrast to the composed authentication and error correction, we can construct a secure quantum channel from a noisy insecure channel and secret key without the need to double encode our qubits in two error-correcting codes.

Theorem 4.4.3. *Let $\#_E$ be a filter introducing the noise given by the depolarizing channel with channel error p. Let $[[n, 1, d]]$ be a family of error-correcting codes with threshold $p_{th} > p$ and decay α. Let π_{AB}^{thr} be the threshold authentication scheme built from these error-correcting codes and with* *threshold parameter* $\tau \in \left(\frac{4p}{3}, 2p_{th}\right)$ *as in Protocol 6. Then for* π_{AB}^{thr} *to obtain* ε*-correctness and* δ*-security, i.e.*

$$\hat{\mathcal{C}}_{\#}||\mathcal{K} \xrightarrow{\pi^{thr},(\varepsilon,\delta)} \mathcal{S}_{\Diamond}\,, \tag{4.58}$$

it is sufficient for the total amount of qubits n_{total} *to grow as*

$$n_{total} \geq \Omega\Big(\max\Big\{\log(1/\varepsilon)^{1/\alpha}, \log(1/\delta)^{1/\alpha}\Big\}\Big)\,. \tag{4.59}$$

Proof. In order to obtain ε-correctness, from Proposition 4.4.1 we need

$$\kappa 2^{-\beta n^{\alpha}\log(p_{\text{th}}/p)} + 2^{-n(\tau-4p/3)^2\log(e)} \leq \varepsilon, \tag{4.60}$$

thus would be enough if both following conditions held,

$$\begin{cases} \kappa 2^{-\beta n^{\alpha}\log(p_{\text{th}}/p)} \leq \varepsilon/2 \\ 2^{-n(\tau-4p/3)^2\log(e)} \leq \varepsilon/2\,. \end{cases} \tag{4.61}$$

For the first item in eq. (4.61), by taking logarithms we need

$$n^{\alpha} \geq \frac{1}{\beta\log(p_{\text{th}}/p)}\log\left(\frac{2\kappa}{\varepsilon}\right), \tag{4.62}$$

and for the second item in eq. (4.61), also by taking logarithms

$$n \geq \frac{1}{(\tau-4p/3)^2\log(e)}\log(2/\varepsilon)\,. \tag{4.63}$$

Since the second condition is weaker, both conditions in eq. (4.61) will hold whenever

$$n \geq \Omega(\log(1/\varepsilon)^{1/\alpha})\,. \tag{4.64}$$

For δ-security, from Proposition 4.4.2 we need

$$\max\left\{\frac{10\kappa}{9\sqrt{2\pi}\sqrt{3b(1-b)}}\exp\left(-\frac{\ln(n)}{2} - \beta n^{\alpha}\ln(p_{th}/b)\right),\right.$$
$$\left.\exp\left(-n\frac{3b}{2}\left(1-\frac{\tau}{2b}\right)^2\right)\right\} \leq \delta\,. \tag{4.65}$$

Note that the above conditions can be simplified as the following two equations holding

$$\begin{cases} A\exp(-\ln(n)/2)\exp(-Bn^{\alpha}) \leq \delta \\ \exp(-nC) \leq \delta\,, \end{cases} \tag{4.66}$$

for the constants

$$A := \frac{10\kappa}{9\sqrt{2\pi}\sqrt{3b(1-b)}}, \quad B := \beta \ln(p_{th}/b) \quad \text{and} \quad C := \frac{3b}{2}\left(1 - \frac{\tau}{2b}\right)^2 . \tag{4.67}$$

By taking logarithms, we see that the first item in eq. (4.66) holds whenever

$$\ln(n)/2 + Bn^\alpha \geq \ln(A/\delta), \quad \text{i.e.} \quad n \geq \Omega\Big(\log(1/\delta)^{1/\alpha}\Big) . \tag{4.68}$$

The second item in eq. (4.66) holds whenever

$$n \geq \frac{1}{C}\ln(1/\delta) , \tag{4.69}$$

a weaker condition than eq. (4.68), thus reaching the desired conclusion. □

4.5 Conclusion

In this chapter we studied the combination of authentication and error-correction in a single primitive, with the goal of reducing the qubit overhead of current methods. We saw that the size blowups of authentication and error correction are (slightly-more than) *multiplied* in a naively composed protocol. As an example of the potential of looking at these properties together, we designed the threshold scheme, for which the resource usage is only dependent on the *maximum* blowup of the two functionalities.

We see that in the usual scenario where the security is more important than the error, $\delta = \varepsilon^{1/\kappa}$ as in the third column of table 4.1, the threshold scheme performs better than both the trap and Clifford scheme. Note that our analysis only shows improvement if the outer error-correcting code encodes qubits independently, because collective encoding might not lead to any overhead.

This is just an initiation to the study of better authentication codes for real, noisy, communications channels, and many questions remain open. For example:

- We leave as an open question what is the maximum gain in efficiency of combining these functionalities. For instance, it is an interesting question whether it is possible to make the Clifford code error-robust in a more efficient way, or if the threshold code is the optimal code.

- In this thesis we only considered the notions of information-theoretic security where the integrity of the plaintext is important, and we do not study key recycling. It could be interesting to combine some of these notions – for instance, to construct a computationally-secure scheme for authentication which also functions as error-correcting code in an efficient way.

Chapter 5

Round-Optimal Oblivious Transfer

This chapter is based on the article *Oblivious Transfer from Zero-Knowledge Proofs* [CMS23], and is joint work with Léo Colisson Palais and Florian Speelman.

The preliminaries have been shortened as the security model has already been treated in Section 2.2. The original section [CMS23, Section 4] has been omitted as we do not treat zero-knowledge on quantum states in this thesis, and hence exclude string oblivious transfer results from this thesis. The proofs of the main theorems have been moved from the appendix to the main body. A conclusion section has been added and the remaining sections have been edited for style and typographical consistency.

In this chapter we give a protocol for round-optimal quantum oblivious transfer (OT). In Section 5.3 we give our OT protocol and prove its security. Note that our OT protocol's achievable security is based on the underlying zero-knowledge (ZK) protocol's. In particular, to achieve round-optimality we need non-interactive ZK (NIZK), and we prove in Section 5.4 that Unruh's NIZK protocol can be made secure in the stand-alone security model with a quantum random oracle.

5.1 Introduction

Oblivious transfer (OT) is a very strong cryptographic primitive, it allows a sender to transfer one of potentially many secrets to a receiver, while remaining oblivious as to which piece has been transferred. In the seminal work by Kilian [Kil88] OT was shown to be enough to perform multi-party computing (MPC).

It is therefore no surprise that the construction of OT has generated widespread interest across the cryptographic community. On the one hand, Impagliazzo [Imp95] showed that all classical OT protocols need to use some structured computational assumptions providing trapdoors, that is, the existence of public-key cryptography is necessary for the existence of OT. On the other hand, once interacting parties have access to quantum computation, Damgård, Fehr, Salvail and Schaffner [DFSS05] showed that OT is possible in the bounded storage model. More recently, Grilo, Lin, Song and Vaikuntanathan [GLSV21] and concurrently Bartusek, Coladangelo, Khurana and Ma [BCKM21] showed that OT is also possible with much weaker assumptions than public-key, based only on (unstructured) one-way functions. There are many reasons to avoid using trapdoor functions. First of all, many of their classical instantiations (based for example on RSA, quadratic residue, elliptic curves) are vulnerable to quantum adversaries. Secondly, it is not yet clear how realistic post-quantum assumptions such as *Learning With Errors* (LWE) are. Therefore, minimizing computational assumptions is an important safetyguard against potential future attacks on the computational assumptions.

Another possible vulnerability for real-life implementations of cryptographic primitives is interaction between parties; the more rounds of communications a task requires, the easier it becomes to compromise, as the adversary can impersonate one of the parties at multiple points of the protocol.

While Peikert, Vaikuntanathan and Waters [PVW08] and concurrently Brakerski and Döttling [BD18] showed (even classical) 2-message OT protocols, optimal in term of round complexity, are achievable using trapdoors, there is no known round-optimal protocol requiring no structure. Therefore, in this chapter we will study the following question:

Does there exist two-message quantum chosen-input bit OT? Based on which computational assumptions?

Our contributions In this chapter, we positively answer the first question by giving a 2-message quantum OT protocol, based on different possible computational assumptions. Recall that the most basic form of OT is a two-party system where Berain (the sender) holds two messages m_0, m_1 and Anboto (the receiver) chooses one of them $b \in \{0, 1\}$, such that Berain must send Anboto *only* message m_b, whilst being oblivious to which b Anboto had choosen.

The main idea of our proof stems from the following naive OT protocol: imagine that after sampling $b \leftarrow \{0, 1\}$ Anboto could prepare two quantum states

$$|\psi^{(b)}\rangle := H|r\rangle \quad \text{and} \quad |\psi^{(1-b)}\rangle := |s\rangle \,, \tag{5.1}$$

for some random bits $r, s \in \{0, 1\}$. She would send $|\psi^{(0)}\rangle$ and $|\psi^{(1)}\rangle$ to Berain, who could rotate the i-th qubit according to the Z^{m_i}-gate and measure both of them in the Hadamard basis, getting the outcomes $z^{(i)}$ that he would send back to Anboto. By the properties of the Z-gate, given a computational basis measurement $\mathcal{M}_{\text{comp}}$, it is easy to see that $z^{(b)} = \mathcal{M}_{\text{comp}}(HH|r \oplus m_b\rangle) = r \oplus m_b$, while $z^{(1-b)} = \mathcal{M}_{\text{comp}}(H|s\rangle)$ is a random bit uncorrelated with m_{1-b}, thus Anboto would only learn the value of the bit m_b. Unfortunately, this protocol is not secure, as Anboto could trivially send two $|+\rangle$ states, enabling her to learn both m_0 and m_1. Starting from this idea, in this chapter we find a way to avoid this cheating strategy by increasing the size of the input states enough so that Berain can check that at least one of the received states is in the computational basis, without having to know exactly the position of this qubit, and without destroying the state.

Theorem 5.1.1 (informal). *There exists a (non-black-box[1]) 2-message bit OT quantum protocol composably secure in the quantum random oracle model, assuming the existence of a collision-resistant second-bit hardcore function.*

At the heart of our approach lies a *Zero-Knowledge* (ZK) type of result on quantum states; meaning that we are able to prove a statement on a quantum state (in this case, being a computational basis state) without revealing anything on that state except the fact that the statement is true. At

[1] Our protocol requires the use of a hash function h: since we need to prove statements on preimages of h in a ZK protocol, this makes our protocol non-black-box with respect to h since the circuit of h must be known to the verifier. Therefore, even if the assumptions on h (collision-resistant and second-bit hardcore) are trivially true if h is modelled as a random oracle, we cannot directly run the ZK protocol on an oracle since the source code of h cannot efficiently be sent to the verifier. For this reason, we do not model h itself as an oracle, and only assume that h is collision-resistant and hiding.

a high level, we define the encoded state as a superposition of pre-images of multiple publicly known images of a given hash function h; and we use classical ZK to prove that the sender knows these, where some of them are tagged as *dummy*. This way, if we prove that one of the two states admits only a single non-dummy preimage (without revealing which state), this state cannot be in superposition of multiple elements, or it would be possible to extract a collision of the hash function. Of course, this assumes that the receiver performs some checks to ensure that the quantum state is a valid encoding, i.e. only contains non-dummy preimages of h: This can be done for instance by checking in superposition that all elements are non-dummy, and by computing h and checking that it belongs to the set of allowed images.

Since we have ZK with different flavours, this actually allow us to obtain a variety of quantum OT protocols, depending on which setup model we want to consider.

i. Quantum random oracle model (QROM): Given the non-interactive zero-knowledge protocol of Unruh [Unr15], we obtain a 2-message bit OT in the QROM.

ii. Plain model: Given the plain model zero-knowledge realization of Hallgren, Smith and Song [HSS11], we obtain 3-message bit OT in the plain model assuming the hardness of LWE.

iii. Common reference string (CRS) model: By Lemma 5.2.6, we obtain 2-message bit OT in the CRS model.

Actually, we prove a much more generic construction, as the OT protocol inherits most of the properties of the underlying ZK protocol.

Theorem 5.1.2 (informal)**.** *Assuming the existence of a collision-resistant hiding one-way function, given a 1-message ZK proof (or argument) of knowledge, we can obtain a 2-message OT protocol in the CRS or RO models, and a 3-message OT protocol in the same security model as the ZK.*

Moreover, if the ZK protocol is secure against any unbounded verifier (resp. prover) and if the function is injective, the resulting OT protocol is secure against any unbounded sender (resp. receiver).

Relation to the literature. Since the introduction of 2-party computing in the seminal article of Yao [Yao82], followed by the famous generalisation to arbitrary many parties of Goldreich, Micali and Wigderson [GMW87], oblivious transfer and multi-party computation have received a tremendous amount of attention [Wie83, PVW08, Rab05, EGL82, CGS02, DGJ$^+$20, KP17, LT22, YAVV22]. The original proposal for quantum OT, by Crépeau and Kilian [CK88], requires 7 messages, which was lowered to 3 messages in the random oracle model by Agarwal, Bartusek, Khurana and Kumar [ABKK23]. There have also been some proposals [PVW08, BD18, Qua20] based on post-quantum assumptions like the LWE problem. Notably Bartusek and Khurana [BK23], after the publication of our article, gave a 1-message OT protocol assuming the hardness of sub-exponential LWE (requiring public-key cryptography), and a 2-message OT in the random oracle setting if the parties have access to a shared EPR pair.

Classical ZK is also a widely studied primitive as it turns out to be extremely useful in many applications; including in MPC, authentication and blockchain protocols [ELE]. Many candidates have been proposed to achieve various ZK flavours: statistical security against malicious prover or malicious verifier, non-interactive or constant round protocols, security in the plain model, CRS, or random oracle [GMR85, Lin13, Unr15, PVW08, BD18, HSS11, PS19, Wat09, AL20, Unr12, BS20a, LMS22]. In this paper, we notably consider the non-interactive ZK protocol of Unruh [Unr15], proven secure in the quantum random oracle model, together with the ZK protocol of Hallgren, Smith and Song [HSS11], proven secure in the plain model assuming the hardness of LWE.

As mentioned in the introduction, our OT protocol inherits the properties of the underlying ZK protocol, thus we include a summary of our results in table 5.1.

5.2 Preliminaries

Although the general quantum-information notation can be found in Chapter 2, here we will introduce some concepts and prior results that we need only in this chapter. In particular, we describe the cryptographic assumptions necessary for our OT protocol, and describe the ideal functionalities for OT and ZK.

5.2.1 Assumptions

Recall how in the introductory chapter to security Section 2.2 we discussed hybrid models inside the stand-alone security model. In this chapter we will work both with the *plain model*, where no assumption on the setup is made,

Article	Setup	Messages	Assumptions	Statistical
[CK88]	Depends	7	OWF	Either
[GLSV21]	Plain model	≥ 7	OWF	No
[BCKM21]	Plain model	≥ 7	OWF	Sender
[ABKK23]	RO	3	OWF	No
[BKS23]	RO + EPR	2	OWF	Yes
This thesis +				
[Unr15]	QROM	2	OWF	No
[HSS11]	Plain model	> 2	LWE	No
S-NIZK	Like ZK	2	Like ZK	Sender
NIZK proof	Like ZK	2	Like ZK	Receiver
ZK	Like ZK	ZK+1 or 2	Like ZK	Like ZK

Table 5.1: Comparison with related works. The party in the *Statistical* column represents the malicious party allowed to be unbounded to get statistical security: Note that using [WW06] we can get statistical security against the other party at the cost of an additional message (of course by losing statistical security against the first party [Lo97]).

the *Common Reference String* (CRS) model and the *Quantum Random Oracle Model* (QROM). We will describe all of them shortly.

The CRS model assumes that a string, honestly sampled according to a fixed procedure, can be shared among all parties. This step is typically not counted in the communication, as in practice we can often heuristically take a publicly known string instead. The ideal functionality modelling the CRS can be described as follows, depicted in figure 5.1.

Definition 5.2.1 (Ideal CRS functionality)**.** *Let* Gen *be a* PPT *sampling procedure. Then the ideal functionality* $\mathcal{F}_{\mathsf{CRS}}^{\mathsf{Gen}}$ *samples a string* $x \leftarrow \mathsf{Gen}(1^\lambda)$ *and outputs the same string* x *to all parties.*

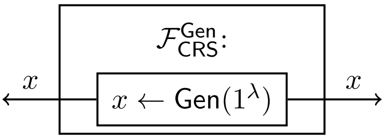
$\mathcal{F}_{\mathsf{CRS}}^{\mathsf{Gen}}$:
x
$x \leftarrow \mathsf{Gen}(1^\lambda)$
x

Figure 5.1: Ideal CRS functionality.

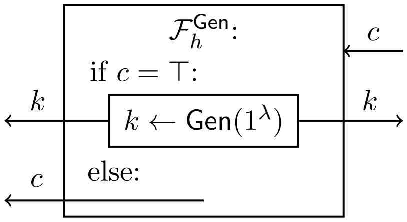
$\mathcal{F}_{h}^{\mathsf{Gen}}$:
c
if $c = \top$:
k
$k \leftarrow \mathsf{Gen}(1^\lambda)$
k
c
else:

Figure 5.2: Ideal collision resistance functionality.

We also require the existence of a particular type of functions for our security proof, which we describe below.

Definition 5.2.2 (Hardcore second-bit)**.** *Let n and m be two polynomials. We say that a function $h\colon \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}$ has a computational (resp. statistical) hardcore second-bit property if for any $s \in \{0,1\}$, any* QPT *(resp. unbounded) adversary $\mathcal{A}$ and for any advice $\{\sigma_\lambda\}_{\lambda\in\mathbb{N}}$:*

$$\left| \Pr_{x \leftarrow \{s\}\times\{0\}\times\{0,1\}^{n(\lambda)}} \Big(\mathcal{A}(1^\lambda, \sigma_\lambda, h(x)) = 1\Big) \right.$$
$$\left. - \Pr_{x \leftarrow \{s\}\times\{1\}\times\{0,1\}^{n(\lambda)}} \Big(\mathcal{A}(1^\lambda, \sigma_\lambda, h(x)) = 1\Big) \right| \leq \mathsf{negl}(\lambda)\,.$$

We extend the definition to a family $\{h_k\colon \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}\}_{k\in\mathcal{K}}$ if for any $k \in \mathcal{K}$, h_k has a computational hardcore second-bit property, and if one can efficiently check for any k whether $k \in \mathcal{K}$ or not.

We note that many functions are expected to have a hardcore second-bit property, since it can be seen as a special case of hiding. This is the case of example of practical hash functions, as they are believed to be hiding.

Definition 5.2.3 (Collision resistance)**.** *Let n and m be two polynomials. A family of functions $\{h_k\colon \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}\}_{k\in\mathcal{K}}$ is said to be (computationally) collision-resistant if there exists a polynomial generation algorithm $k \leftarrow \mathsf{Gen}_h(1^\lambda)$ such that for any $k \in \mathcal{K}$, h_k can be classically evaluated in polynomial time, and for any (potentially non-uniform)* QPT *adversary $\mathcal{A}$ and advice $\{\sigma_\lambda\}_{\lambda\in\mathbb{N}}$:*

$$\Pr_{\substack{k \leftarrow \mathsf{Gen}_h(1^\lambda),\\ (x,x') \leftarrow \mathcal{A}(k,\sigma_\lambda)}} \big(x \neq x' \wedge h_k(x) = h_k(x')\big) \leq \mathsf{negl}(\lambda)\,.$$

Note that even if we heuristically expect the protocol to stay secure when we replace h_k with a fixed hash function like SHA-256, to prove the security we need to sample the function h_k after the beginning of the protocol. The reason being that the adversaries are non-uniform, and the advice could contain a collision if it was chosen after h_k.

We define the functionality that distributes the value of h_k as follows, depicted in figure 5.2.

Definition 5.2.4. *Let $\{h_k \colon \{0,1\}^{l(\lambda)} \to \{0,1\}^{m(\lambda)}\}_{k\in\mathcal{K}}$ be a family of collision resistant functions generated by Gen, with a hardcore second-bit property. We define the ideal functionality $\mathcal{F}_h^{\mathsf{Gen}}$ as follows. $\mathcal{F}_h^{\mathsf{Gen}}$ receives an input c from Berain's interface, if $c = \top$, the functionality samples $k \leftarrow \mathsf{Gen}(1^\lambda)$ and sends k to both parties, otherwise if $c \in \mathcal{K}$, it forwards c to Anboto's interface. The ideal party $\mathbf{A}_I$ just forwards the received k, while the ideal party $\mathbf{B}_I$ sends $c = \top$ to the functionality and outputs the received k.*

The above functionality can be realized either in the plain-model or in the CRS model, attending to the party that is tasked with sampling h_k. On the one hand, we could let Berain honestly sample h_k, which can be done in the plain-model at the cost of one message from Berain to Anboto.

Lemma 5.2.5 ($\mathcal{F}_h^{\mathsf{Gen}}$ in the plain model)**.** *The 1-message protocol where Berain samples $x \leftarrow \mathsf{Gen}(1^\lambda)$ and sends x to Anboto, and Anboto outputs x only if $x \in \mathcal{K}$ realizes the functionality $\mathcal{F}_h^{\mathsf{Gen}}$ in the plain model.*

Proof. If no party is corrupted (correctness), then the outputs of both parties are always distributed according to $\mathcal{F}_{\mathsf{CRS}}^{\mathsf{Gen}}$, so no environment can even statistically distinguish between the ideal and real world. If Anboto gets corrupted, we define the simulator exactly as the malicious Anboto $\hat{\mathbf{A}}$, and both worlds are identical (up to the name we give to the different parts of the world) and therefore indistinguishable. If Berain $\hat{\mathbf{B}}$ is malicious, then we define the simulator that runs $\hat{\mathbf{B}}$ and forwards the output to the ideal functionality: the test performed by the functionality is exactly the test performed by Anboto in the real world, so both worlds are equal. □

We could instead assume that the circuit of h_k is provided by a CRS, which requires no additional round of communication, but we are not anymore in the plain model.

Lemma 5.2.6 ($\mathcal{F}_h^{\mathsf{Gen}}$ in the CRS model)**.** *In the $\mathcal{F}_{\mathsf{CRS}}^{\mathsf{Gen}}$-hybrid model, the trivial 0-message protocol where both Anboto and Berain output the value given by $\mathcal{F}_{\mathsf{CRS}}^{\mathsf{Gen}}$ realizes the $\mathcal{F}_h^{\mathsf{Gen}}$ functionality.*

Proof. The proof is trivial and very similar to Lemma 5.2.5. The only case where it differs is when Berain is malicious. In this case, we define the simulator as outputting $c = \top$ to the functionality, and then feeding the received advice and k to $\hat{\mathbf{B}}$. Again, the ideal and real worlds are equal, which concludes the proof. □

Finally, we can reduce the round complexity of our protocols by assuming the quantum random oracle model.

Definition 5.2.7. *The ideal functionality $\mathcal{F}_H$ for a function $H\colon \{0,1\}^* \to \{0,1\}^n$, samples on any classical input $x \in \{0,1\}^*$ a random output $y \in \{0,1\}^n$. For general quantum queries $\sum_x \alpha_x |x\rangle|y\rangle$, the functionality $\mathcal{F}_H$ applies the unitary*

$$U_H\left(\sum_x \alpha_x |x\rangle|y\rangle\right) := \sum_x \alpha_x |x\rangle|y \oplus H(x)\rangle\,.$$

We also assume a distribution ROdist on functions, modelling the distributions of a random oracle, that is, ROdist would be the uniform distribution on $\{0,1\}^* \to \{0,1\}^n$.

5.2.2 Definitions

As explained in Section 2.2, we prove security of a protocol in the stand-alone model by giving an ideal functionality. We include here the ideal functionalities we will be working with for the rest of the chapter. The definitions are taken from [HSS11].

Definition 5.2.8 (Ideal OT functionality)**.** *We define the ideal functionality $\mathcal{F}_{OT}$ for Oblivious Transfer (OT) as follows. The ideal functionality $\mathcal{F}_{OT}$ receives two messages m_0 and m_1 from Berain's interface (or an abort message) and one bit $b \in \{0,1\}$ from Anboto's interface (or an abort message), if no party decided to abort, it sends m_b to Anboto. We define trivially the dummy parties $\tilde{\Pi} = (\tilde{\mathbf{A}}, \tilde{\mathbf{B}})$ that forward the inputs/outputs to/from $\mathcal{F}_{OT}$.*

For a pictorial representation of $\mathcal{F}_{OT}$ see figure 5.3.

Classical zero-knowledge (ZK) proofs allow a *prover* to prove a statement to the *verifier* without revealing anything beyond the fact that this statement is true. Our protocol for OT use a ZK protocol as a blackbox. We say that a relation $\mathcal{R}$ describes a language $\mathcal{L}$ whenever an *instance* $x \in \mathcal{L}$ if and only if there exists a *witness* w such that $x\mathcal{R}w$.

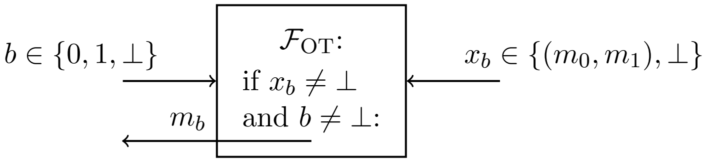

Figure 5.3: Ideal OT functionality.

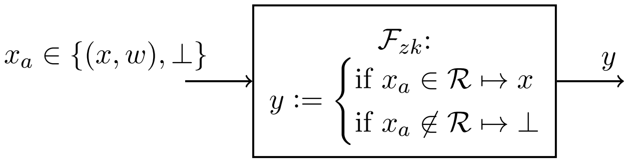
$x_a \in \{(x, w), \perp\}$
$\mathcal{F}_{zk}$:
$y := \begin{cases} \text{if } x_a \in \mathcal{R} \mapsto x \\ \text{if } x_a \notin \mathcal{R} \mapsto \perp \end{cases}$
y

Figure 5.4: Ideal ZK functionality.

Definition 5.2.9 (Ideal ZK functionality)**.** *We define the ideal functionality $\mathcal{F}_{zk}$ for Zero-Knowledge (ZK) for the relation $\mathcal{R}$ describing the language $\mathcal{L}$ as follows. The ideal functionality $\mathcal{F}_{zk}$ receives (x, w) from Anboto's interface (or an abort message), if $x\mathcal{R}w$ then the functionality sends x to Berain, and otherwise sends an abort to Berain's interface instead.*

In ZK it is customary to call Anboto a *prover*, denoted by P and Berain a *verifier*, denoted V. For a pictorial representation of $\mathcal{F}_{OT}$ see figure 5.4. Note that the ZK functionality also implies that the ZK protocol is a *proof of knowledge* (PoK) protocol as the functionality can extract the witness (quantumly this is also known as *state-preserving*, as extracting the witness should not disturb the state of the adversary).

Moreover, we often distinguish between ZK protocols attending to the malicious parties security is proven against. In the quantum stand-alone formalism, recall Section 2.2, ZK *proofs* are protocols that $\mathsf{C-QSA}$ realize $\mathcal{F}_{zk}^{\mathcal{R}}$ secure against a malicious prover and *statistical* ZK are protocols that $\mathsf{C-QSA}$ realize $\mathcal{F}_{zk}^{\mathcal{R}}$ secure against a malicious verifier.

There are multiple protocols realising the ZK functionality $\mathcal{F}_{zk}^{\mathcal{R}}$, either in the plain model [HSS11] or non-interactively in the QROM [Unr15]. This last work is not expressed in the quantum stand-alone model, but we prove in Section 5.4 that it can be reformulated in this framework.

5.2.3 Useful lemmas

Lemma 5.2.10. *Let $|\phi\rangle$ and $|\phi'\rangle$ be two orthogonal pure states and $\beta \in [0,1]$. We consider the normalized state $|\psi\rangle = \sqrt{1-\beta}|\phi\rangle + \sqrt{\beta}|\phi'\rangle$. Then, $\mathrm{TD}(|\phi\rangle, |\psi\rangle) = \sqrt{\beta}$.*

Proof. This is a direct consequence of the trace distance between two pure states:

$$\mathrm{TD}(|\phi\rangle, |\psi\rangle) = \sqrt{1 - |\langle\phi|\psi\rangle|^2} = \sqrt{1 - |\langle\phi|(\sqrt{1-\beta}|\phi\rangle + \sqrt{\beta}|\phi'\rangle)|^2} \quad (5.2)$$

$$= \sqrt{1 - |\sqrt{1-\beta}|^2} = \sqrt{\beta}\,. \quad (5.3)$$

□

5.3 Our protocol for bit oblivious transfer

In this section we present our oblivious transfer (OT) protocol and prove that it computationally quantum stand-alone ($\mathsf{C-QSA}$) realizes the ideal OT functionality $\mathcal{F}_{OT}$.

We present only one proof, which allows for multiple instantiations. First, we require collision resistant hash functions with the hardcore second-bit property $\mathcal{F}_h^{\mathsf{Gen}}$, which can be realized either in the plain model with 1 message from Berain to Anboto (see Lemma 5.2.5) or non-interactively in the CRS model (see Lemma 5.2.6). Second, we require classical zero-knowledge $\mathcal{F}_{zk}^{\mathcal{R}}$, whose properties get lifted to our OT protocol. In particular, there exists a non-interactive zero-knowledge realization in the QROM, which we prove in the next section.

Protocol 7 Protocol for (possibly 2-message) chosen bit Oblivious Transfer.

Inputs. Anboto gets $b \in \{0,1\}$ and Berain gets $m_0, m_1 \in \{0,1\}$.

Assumption. We assume Π_{zk} is a (NI)ZK protocol for $\mathcal{F}_{zk}^{\mathcal{R}}$, and h is a collision-resistant second-bit hardcore function $\mathcal{F}_h^{\mathsf{Gen}}$ distributed by Π_h.

Goal. Anboto to learn *only* Berain's message m_b, without revealing b.

The protocol:

1: Anboto samples $r, s \leftarrow \{0,1\}$, two witnesses for b:

$$\omega_0^{(b)}, \omega_1^{(b)} \leftarrow \{0\} \times \{0,1\}^n\,,$$

and two witnesses for $1-b$:

$$\omega_s^{(1-b)} \leftarrow \{0\} \times \{0,1\}^n \quad \text{and} \quad \omega_{1-s}^{(1-b)} \leftarrow \{1\} \otimes \{0,1\}^n \,.$$

2: Anboto computes the four hash values $h_d^{(c)} := h(d||\omega_d^{(c)})$, with $c, d \in \{0,1\}$.

3: Anboto and Berain run the (NI)ZK protocol Π_{zk} to prove that there exist witnesses $\{\omega_d^{(c)}\}_{c,d\in\{0,1\}}$ for the hashed values $h_d^{(c)} = h(d||\omega_d^{(c)})$ such that at least one for them has a 1 at the first bit, i.e. $\omega_d^{(c)}[1] = 1$ for some $c, d \in \{0,1\}$.

4: Anboto sends Berain the (NI)ZK proof and the hash values $\{h_d^{(c)}\}_{c,d\in\{0,1\}}$, together with the quantum states

$$|\psi^{(b)}\rangle := |0\rangle|\omega_0^{(b)}\rangle + (-1)^r|1\rangle|\omega_1^{(b)}\rangle \quad \text{and} \quad |\psi^{(1-b)}\rangle := |s\rangle|\omega_s^{(1-b)}\rangle \,.$$

5: Berain verifies the (NI)ZK proof. He then verifies that the quantum states are honestly prepared by applying on $|\psi^{(c)}\rangle|0\rangle$ the following unitary:

$$x, \omega \mapsto \omega[1] \neq 1 \wedge h(x||\omega) \in \{h_0^{(c)}, h_1^{(c)}\}\,,$$

and measuring the auxiliary register to check if it is equal to 1. If not, he aborts the protocol.

6: Berain measures the second registers of $|\psi^{(0)}\rangle$ and $|\psi^{(1)}\rangle$ in the Hadamard basis to obtain outcomes $w^{(0)}$ and $w^{(1)}$ respectively. Note that at this step, the first register of $|\psi^{(b)}\rangle$ contains a $|\pm\rangle$ state while $|\psi^{(1-b)}\rangle$ contains $|s\rangle$, without Berain knowing b.

7: Berain applies for each $c \in \{0,1\}$ the Z^{m_c}-gate on $|\psi^{(c)}\rangle$ and measures it in the Hadamard basis to obtain outcomes $z^{(c)}$.

8: Berain sends the classical outcomes $\{w^{(c)}, z^{(c)}\}_{c\in\{0,1\}}$ to Anboto.

9: Anboto computes $r \oplus \langle w^{(b)}, \omega_0^{(b)} \oplus \omega_1^{(b)}\rangle \oplus z^{(b)} = m_b$.

Theorem 5.3.1. *Let $\Pi_h = (\mathbf{A}_h, \mathbf{B}_h)$ be a protocol that $\mathsf{C-QSA}$ realizes $\mathcal{F}_h^{\mathsf{Gen}}$ for a family $\{h_k\}_{k\in\mathcal{K}}$ of collision resistant hash functions with the hardcore second-bit property. Let $\Pi_{zk} = (\mathbf{A}_{zk}, \mathbf{B}_{zk})$ be a protocol that $\mathsf{C-QSA}$ realizes the (NI)ZK functionality $\mathcal{F}_{zk}^{\mathcal{R}}$, where $(h_0^0, h_0^1, h_1^0, h_1^1)\mathcal{R}(\omega_0^0, \omega_0^1, \omega_1^0, \omega_1^1) \Leftrightarrow \forall c, d, h(d||\omega_d^{(c)}) = h_d^{(c)}$. Then Protocol 7, where h is obtained by first running Π_h, $\mathsf{C-QSA}$ realizes the OT functionality $\mathcal{F}_{OT}$. In particular,*

i. The realization is secure against malicious Berain if and only if h has the statistical hardcore second-bit property.

ii. The realization is secure against malicious Anboto if and only if for any $k \in \mathcal{K}$, h_k is injective.

Proof. We will split the proof depending on the parties that the static adversary $\mathcal{A}$ can corrupt, that is $\mathcal{A} \in \{\emptyset, \mathbf{A}, \mathbf{B}\}$.

Case 1: correctness (no corrupted party). If no party is corrupted, $\mathcal{A} = \emptyset$, we are proving that the protocol is correct. First, we can cut $\mathbf{A}$ and $\mathbf{B}$ in four parts: the part that runs $\mathbf{A}_h$, the part that generates $h_c^{(c)}$, the part that runs the ZK proof, and the rest. Because of the correctness of the protocol distributing h, and because of the completeness of the ZK protocol, we can indistinguishably replace the first and third parts of $\mathbf{A}$ and $\mathbf{B}$ with the ideal dummy parties and the corresponding functionalities. Now, because $\omega_{1-l}^{(1-b)}[1] = 1$, the statement that we prove in the (NI)ZK proof is true, so by the completeness of the (NI)ZK protocol the check succeeds. Then, during the second step we apply the unitary that maps $|x\rangle|\omega\rangle|0\rangle$ to $|x\rangle|\omega\rangle|1\rangle$ only if $\omega[1] \neq 1$ and if $h(x||\omega)$ appears in the list of hashes, which is true for all terms appearing in $|\psi^{(b)}\rangle$ and $|\psi^{(1-b)}\rangle$ by construction, so after this step the two states become:

$$|\psi^{(b)}\rangle \rightsquigarrow |0\rangle|\omega_0^{(b)}\rangle|1\rangle + (-1)^r|1\rangle|\omega_1^{(b)}\rangle|1\rangle = |\psi^{(b)}\rangle|1\rangle\,, \tag{5.4}$$

$$|\psi^{(1-b)}\rangle \rightsquigarrow |s\rangle|\omega_s^{(b)}\rangle|1\rangle = |\psi^{(1-b)}\rangle|1\rangle\,. \tag{5.5}$$

Because the last register of each state is not entangled with their respective first two registers, measuring them will output 1 in both cases and will not disturb the states on the first two registers. Then, we measure the second register of each state in the Hadamard basis, i.e. we first apply Hadamard gates on all qubits and then we measure in the computational basis.

After the Hadamard gates, the state $|\psi^{(b)}\rangle$ is turned into (omitting the constants):

$$(I \otimes H^{n+1})|\psi^{(b)}\rangle \tag{5.6}$$

$$= |0\rangle H^{n+1}|\omega_0^{(b)}\rangle + (-1)^r|1\rangle H^{n+1}|\omega_1^{(b)}\rangle \tag{5.7}$$

$$\begin{aligned} &= |0\rangle \sum_{w^{(b)} \in \{0,1\}^{n+1}} (-1)^{\langle w^{(b)}, \omega_0^{(b)}\rangle}|w^{(b)}\rangle \\ &\quad + (-1)^r|1\rangle \sum_{w^{(b)} \in \{0,1\}^{n+1}} (-1)^{\langle w^{(b)}, \omega_1^{(b)}\rangle}|w^{(b)}\rangle \end{aligned} \tag{5.8}$$

$$= |0\rangle + |1\rangle \left(\sum_{w^{(b)} \in \{0,1\}^{n+1}} (-1)^{r \oplus \langle w^{(b)}, \omega_0^{(b)} \oplus \omega_1^{(b)} \rangle} |w^{(b)}\rangle \right) . \tag{5.9}$$

If we measure then an outcome $w^{(b)}$ and define $\alpha := r \oplus \langle w^{(b)}, \omega_0^{(b)} \oplus \omega_1^{(b)} \rangle$, the above state collapses to the Hadamard basis state:

$$|\phi^{(b)}\rangle := |0\rangle + (-1)^{\alpha} |1\rangle . \tag{5.10}$$

For the second state $|\psi^{(1-b)}\rangle = |s\rangle|\omega_s^{(1-b)}\rangle$, the first register is not entangled with the second, so measuring the second register does not disturb the first one. Thus we end up with the computational basis state:

$$|\phi^{(1-b)}\rangle := |s\rangle . \tag{5.11}$$

In the final step, Berain rotates the two states

$$Z^{m_b} |\phi^{(b)}\rangle = |0\rangle + (-1)^{\alpha \oplus m_b} |1\rangle \quad \text{and} \quad Z^{m_{1-b}} |\phi^{(1-b)}\rangle = |s\rangle . \tag{5.12}$$

Note that the rotation does nothing in the second case, hence all information about m_{1-b} is lost. When Berain measures both states in the Hadamard basis, the first outcome is $z^{(b)} = \alpha \oplus m_b$, and the second outcome $z^{(1-b)}$ is just a random bit. At the end of the protocol Anboto outputs $\alpha \oplus z^{(b)} = m_b$, which ends the proof of correctness.

Case 2: malicious sender (Berain). We consider now the case where the adversary $\mathcal{A}$ corrupts the sender Berain, that is $\mathcal{A} = \hat{\mathbf{B}}$.

First, we notice that since the functionality $\mathcal{F}_{zk}^{\mathcal{R}}$ sends the word x (in our case the hashes $\{h_d^{(c)}\}_{c,d \in \{0,1\}}$) to the verifier, we do not need to send x another time before (if needed we can assume that the (NI)ZK protocol starts by sending x). Let $\hat{\mathbf{B}}$ be an adversary. Without loss of generality[2], we can decompose $\hat{\mathbf{B}}$ into $\hat{\mathbf{B}}_0$ and $\hat{\mathbf{B}}_1$, where $\hat{\mathbf{B}}_0$ is the QIM running during the (NI)ZK protocol (so it receives an arbitrary advice σ and interacts with $\mathbf{A}$ during the (NI)ZK protocol), forwarding its final internal state to $\hat{\mathbf{B}}_1$ that runs the rest of the protocol (in particular $\hat{\mathbf{B}}_1$ receives the quantum state and is supposed to output some measurement outcomes). For any environment $(\mathbf{Z}, \sigma)$ we define below the following hybrids (see Protocols 8 to 14 for algorithmic descriptions):

[2] If the ZK is non-interactive, $\hat{\mathbf{B}}_0$ would just store the proof π, the hashes $(h_{i,j})$ and advice σ and forward it to $\hat{\mathbf{B}}_1$.

i. $\mathsf{World}_0 := \mathsf{REAL}^{\sigma}_{\Pi,\hat{\mathbf{B}},\mathbf{Z}}$ is the real world (where Anboto runs internally the (NI)ZK protocol $\mathbf{A}_{zk}$ with $\hat{\mathbf{B}}$).

ii. World_1 is like World_0 except that the Π_h protocol (in charge of sharing h) is replaced by the simulated version Sim_h.

iii. World_2 is like World_0 except that the (NI)ZK protocol is replaced by the simulated version Sim_{zk}.

iv. World_3 is like World_2 except that we remove $\mathcal{F}^{\mathcal{R}}_{zk}$ and always forward $\{h^c_d\}_{c,d\in\{0,1\}}$ to $\mathsf{Sim}_{\hat{\mathbf{B}}_0}$.

v. World_4 is like World_3 except that we sample $\omega^{(1-b)}_{1-s} \leftarrow \{0\} \times \{0,1\}^n$.

vi. World_5 is like World_4 except that we define instead

$$|\psi^{(1-b)}\rangle := |0\rangle|\omega^{(1-b)}_0\rangle + (-1)^{r'}|1\rangle|\omega^{(1-b)}_1\rangle\,, \tag{5.13}$$

where $r' \leftarrow \{0,1\}$ is a random bit.

vii. World_6 is like World_5 except that we reorder some operations and we cut Anboto in three parts: a simulator $\mathsf{Sim}_{\hat{\mathbf{B}}}$ (the simulator will also absorb $\mathsf{Sim}_{\hat{\mathbf{B}}_0}$ and $\hat{\mathbf{B}}_1$ by simply forwarding the input σ to $\mathsf{Sim}_{\hat{\mathbf{B}}_0}$ and the output of $\hat{\mathbf{B}}_1$ to $\mathbf{Z}$), the ideal functionality $\mathcal{F}_{OT}$, and the dummy party $\tilde{\mathbf{A}}$ that forwards b to $\mathcal{F}_{OT}$ and outputs the answer m_b of $\mathcal{F}_{OT}$ (in Protocol 14 $\mathcal{F}_{OT}$ and $\tilde{\mathbf{A}}$ are drawn together for to save space). More precisely, we see that all the messages sent to $\hat{\mathbf{B}}$ are sampled exactly in the same way, irrespective of the value of b, so we can push that outside of Anboto into the simulator. The only part that still depends on b is the output message. To avoid this dependency, the simulator will compute the two outputs (when $b = 0$ and when $b = 1$) and send them to $\mathcal{F}_{OT}$ that will be in charge of outputting the appropriate value. This way, we see that $\mathsf{World}_6 = \mathsf{IDEAL}^{\sigma,\mathcal{F}_{OT}}_{\tilde{\Pi},\mathsf{Sim}_{\hat{\mathcal{B}}},\mathbf{Z}}$.

First, we see that $\mathsf{World}_0 \approx \mathsf{World}_2$ because we assumed that the underlying protocol $\mathsf{C-QSA}$ realizes $\mathcal{F}^{\mathcal{R}}_{zk}$: If it were not the case, then we could easily break the $\mathsf{C-QSA}$ property of $\mathcal{F}^{\mathcal{R}}_{zk}$ by merging the classical sampling procedure inside σ to get a new σ' (the value of the witness being kept as a side-information for $\mathbf{Z}$ in σ') and the rest of the procedure (preparation of the quantum state and running $\hat{\mathbf{B}}_1$) inside $\mathbf{Z}$ to produce a new $\mathbf{Z}'$ able to attack $\mathcal{F}^{\mathcal{R}}_{zk}$ with exactly the same probability.

Protocol 8 Case 2 (malicious Berain): World_0.

	$\mathbf{A}(b)$		$\hat{\mathbf{B}}(\sigma_{\hat{\mathbf{B}}})$		$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	Run $\mathbf{A}_h$ to obtain h	$\longleftrightarrow$	$\hat{\mathbf{B}}_0$		
2:	$\forall c, d$, sample $\omega_d^{(c)}$ and $h_d^{(c)}$ as in the main protocol				
3:	Run $\mathbf{A}_{zk}(\forall c, d, \omega_d^{(c)}, h_d^{(c)})$	$\longleftrightarrow$	$\hat{\mathbf{B}}_1$		
4:	Sample $\vert\psi^{(0)}\rangle, \vert\psi^{(1)}\rangle$ as in the main protocol	$\xrightarrow{\vert\psi^{(0)}\rangle, \vert\psi^{(1)}\rangle}$			
5:		$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$	$\hat{\mathbf{B}}_2$	$\xrightarrow{\text{state}}$	
6:	Compute α as in the main protocol	$\xrightarrow{\alpha \oplus z^{(b)}}$			**return** anything

Protocol 9 Case 2 (malicious Berain): World_1.

	$\mathbf{A}(b)$		$\hat{\mathbf{B}}(\sigma_{\hat{\mathbf{B}}})$		$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	Run $\mathcal{F}_h^{\mathsf{Gen}}$ to obtain h	$\xleftarrow{h}$	$\mathsf{Sim}_{h,\hat{\mathbf{B}}_0}$		
2:	$\forall c, d$, sample $\omega_d^{(c)}$ and $h_d^{(c)}$ as in the main protocol				
3:	Run $\mathbf{A}_{zk}(\forall c, d, \omega_d^{(c)}, h_d^{(c)})$	$\longleftrightarrow$	$\hat{\mathbf{B}}_1$		
4:	Sample $\vert\psi^{(0)}\rangle, \vert\psi^{(1)}\rangle$ as in the main protocol	$\xrightarrow{\vert\psi^{(0)}\rangle, \vert\psi^{(1)}\rangle}$			
5:		$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$	$\hat{\mathbf{B}}_2$	$\xrightarrow{\text{state}}$	
6:	Compute α as in the main protocol	$\xrightarrow{\alpha \oplus z^{(b)}}$			**return** anything

Protocol 10 Case 2 (malicious Berain): World_2.

	$\mathbf{A}(b)$		$\hat{\mathbf{B}}(\sigma_{\hat{\mathbf{B}}})$		$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	Run $\mathcal{F}_h^{\mathsf{Gen}}$ to obtain h	$\xleftarrow{\quad h \quad}$	$\mathsf{Sim}_{h,\hat{\mathbf{B}}_0}$		
2:	$\forall c, d$, sample $\omega_d^{(c)}$ and $h_d^{(c)}$ as in the main protocol				
3:	Run $\mathcal{F}_{zk}^{\mathcal{R}}$	$\xrightarrow{\forall c, d, h_d^{(c)}}$	$\mathsf{Sim}_{zk,\hat{\mathbf{B}}_1}$		
4:	Sample $\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle$ as in the main protocol	$\xrightarrow{\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle}$			
5:		$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$	$\hat{\mathbf{B}}_2$	$\xrightarrow{\text{state}}$	
6:	Compute α as in the main protocol	$\xrightarrow{\alpha \oplus z^{(b)}}$			**return** anything

Protocol 11 Case 2 (malicious Berain): World_3.

	$\mathbf{A}(b)$		$\hat{\mathbf{B}}(\sigma_{\hat{\mathbf{B}}})$		$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	Run $\mathcal{F}_h^{\mathsf{Gen}}$ to obtain h	$\xleftarrow{\quad h \quad}$	$\mathsf{Sim}_{h,\hat{\mathbf{B}}_0}$		
2:	$\forall c, d$, sample $\omega_d^{(c)}$ and $h_d^{(c)}$ as in the main protocol				
3:	~~Run $\mathcal{F}_{zk}^{\mathcal{R}}$~~	$\xrightarrow{\forall c, d, h_d^{(c)}}$	$\mathsf{Sim}_{zk,\hat{\mathbf{B}}_1}$		
4:	Sample $\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle$ as in the main protocol	$\xrightarrow{\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle}$			
5:		$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$	$\hat{\mathbf{B}}_2$	$\xrightarrow{\text{state}}$	
6:	Compute α as in the main protocol	$\xrightarrow{\alpha \oplus z^{(b)}}$			**return** anything

Protocol 12 Case 2 (malicious Berain): World_4.

	$\mathbf{A}(b)$		$\hat{\mathbf{B}}(\sigma_{\hat{\mathbf{B}}})$		$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	Run $\mathcal{F}_h^{\mathsf{Gen}}$ to obtain h	$\xleftarrow{\quad h \quad}$	$\mathsf{Sim}_{h,\hat{\mathbf{B}}_0}$		
2:	$\forall c, d$, sample $\omega_d^{(c)}$ and $h_d^{(c)}$ as in the main protocol, except for $\omega_{1-s}^{(1-b)} \leftarrow \{0\} \times \{0,1\}^n$	$\xrightarrow{\forall c, d, h_d^{(c)}}$	$\mathsf{Sim}_{zk,\hat{\mathbf{B}}_1}$		
3:	Sample $\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle$ as in the main protocol	$\xrightarrow{\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle}$			
4:		$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$	$\hat{\mathbf{B}}_2$	$\xrightarrow{\text{state}}$	
5:	Compute α as in the main protocol	$\xrightarrow{\alpha \oplus z^{(b)}}$			**return** anything

Then, $\mathsf{World}_2 = \mathsf{World}_3$: by construction, there always exists a witness starting with a 0, so $\mathcal{F}_{zk}^{\mathcal{R}}$ will always forward $\{h_d^c\}_{c,d \in \{0,1\}}$.

We also have $\mathsf{World}_3 \approx \mathsf{World}_4$ because of the hardcore second-bit property (see Definition 5.2.2). Otherwise, we can easily break the hardcore second-bit property by defining $\mathcal{A}$ as World_3 except that the $\omega_{1-s}^{(1-b)}$ is sampled externally by the "challenger" that only provides $h_{1-s}^{(1-b)}$ to $\mathcal{A}$ (note that $\omega_{1-s}^{(1-b)}$ is not needed here except to compute $h_{1-s}^{(1-b)}$). Note that this attacker is only valid if the computational power of $\hat{\mathbf{B}}$ is lower than the computational power defined in the hardcore second-bit property.

$\mathsf{World}_4 = \mathsf{World}_5$, because for any x and x', it is statistically impossible to distinguish $\lvert x\rangle + (-1)^{r'}\lvert x'\rangle$ from $\lvert x''\rangle$ where x'' equals x with probability $\frac{1}{2}$ and x' otherwise, even given x and x' (the random sign r' must however stay hidden). Indeed, we can simply study the density matrices

$$\lvert x\rangle \pm \lvert x'\rangle = \frac{1}{2}\left(\frac{\lvert x\rangle + \lvert x'\rangle}{\sqrt{2}}\frac{\langle x\rvert + \langle x'\rvert}{\sqrt{2}} + \frac{\lvert x\rangle - \lvert x'\rangle}{\sqrt{2}}\frac{\langle x\rvert - \langle x'\rvert}{\sqrt{2}}\right) \tag{5.14}$$

Protocol 13 Case 2 (malicious Berain): World_5.

	$\mathbf{A}(b)$		$\hat{\mathbf{B}}(\sigma_{\hat{\mathbf{B}}})$		$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	Run $\mathcal{F}_h^{\mathsf{Gen}}$ to obtain h	$\xleftarrow{\quad h \quad}$	$\mathsf{Sim}_{h,\hat{\mathbf{B}}_0}$		
2:	$\forall c, d$, sample $\omega_d^{(c)}$ and $h_d^{(c)}$ as in the main protocol, except for $\omega_{1-s}^{(1-b)} \leftarrow \{0\} \times \{0,1\}^n$	$\xrightarrow{\forall c, d, h_d^{(c)}}$	$\mathsf{Sim}_{zk,\hat{\mathbf{B}}_1}$		
3:	Sample $\lvert\psi^{(b)}\rangle$ as in the main protocol, $r' \leftarrow \{0,1\}$ and set $\lvert\psi^{(1-b)}\rangle := \lvert 0\rangle\lvert\omega_0^{(1-b)}\rangle + (-1)^{r'}\lvert 1\rangle\lvert\omega_1^{(1-b)}\rangle$	$\xrightarrow{\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle}$			
4:		$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$	$\hat{\mathbf{B}}_2$	$\xrightarrow{\text{state}}$	
5:	Compute α as in the main protocol	$\xrightarrow{\alpha \oplus z^{(b)}}$			**return** anything

Protocol 14 Case 2 (malicious Berain): World_6.

	$\mathbf{A}_I(b)$ and $\mathcal{F}_{OT}$		$\mathsf{Sim}_{\hat{\mathbf{B}}}(\sigma_{\hat{\mathbf{B}}})$		$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:			$(h, \mathsf{state}_0) \leftarrow \mathsf{Sim}_{h,\hat{\mathbf{B}}_0}$		
2:			Check if $h \in \mathcal{K}$		
3:			$\forall c, d$, sample $\omega_d^{(c)}$ and $h_d^{(c)}$ as in the main protocol, except for $\omega_{1-s}^{(1-b)} \leftarrow \{0\} \times \{0,1\}^n$		
4:			Sample $\lvert\psi^{(b)}\rangle$ as in the main protocol, $r' \leftarrow \{0,1\}$ and set $\lvert\psi^{(1-b)}\rangle := \lvert 0\rangle\lvert\omega_0^{(1-b)}\rangle$ $+(-1)^{r'}\lvert 1\rangle\lvert\omega_1^{(1-b)}\rangle$		
5:			$\forall c(w^{(c)}, z^{(c)}, \mathsf{state}_2)$ $\leftarrow \hat{\mathbf{B}}_2(\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle, \mathsf{state}_1)$	$\xrightarrow{\mathsf{state}_2}$	
6:		$\xleftarrow{m^{(0)}, m^{(1)}}$	$\forall c, m^{(c)} = (r^{(c)} \oplus$ $\bigoplus_i w^{(c)}[i](\omega_0^{(c)} \oplus \omega_1^{(c)})[i]) \oplus z^{(c)}$		
7:	Run $\mathcal{F}_{OT}$ to obtain $m^{(b)}$		$\xrightarrow{m^{(b)}}$		**return** anything

$$= \frac{1}{4}\big(|x\rangle\langle x| + |x\rangle\langle x'| + |x'\rangle\langle x| \\ +|x'\rangle\langle x'| + |x\rangle\langle x| - |x\rangle\langle x'| - |x'\rangle\langle x| + |x'\rangle\langle x'|\big) \tag{5.15}$$

$$= \frac{1}{2}\big(|x\rangle\langle x| + |x'\rangle\langle x'|\big) . \tag{5.16}$$

But this last expression is exactly the density matrix of $|x''\rangle$ where x'' equals x with probability $1/2$ and x' otherwise.

In our case, we have $x = 0\|\omega_0^{(1-b)}$, $x' = 1\|\omega_1^{(1-b)}$ and $x'' = s\|\omega_s^{(1-b)}$ since s plays the role of the random coin determining the element to send in World_4 (you can see that s plays no role in World_5 since we can re-order the steps in the sampling procedure). Since s and r' are only used to determine $|\psi^{(1-b)}\rangle$ and since both worlds are equal beside the choice of $|\psi^{(1-b)}\rangle$, we can conclude that $\mathsf{World}_4 = \mathsf{World}_5$ as otherwise we could factor out $|\psi^{(1-b)}\rangle$ from the worlds and use the remaining part to distinguish between $|x\rangle + (-1)^{r'}|x'\rangle$ from $|x''\rangle$ which is physically impossible as we just saw.

Finally, $\mathsf{World}_5 = \mathsf{World}_6$ since we just reordered the sampling procedure, delocalized their computation to $\mathsf{Sim}_{\hat{\mathbf{B}}}$ and used $\mathcal{F}_{OT}$ to discard the message corresponding to m_{1-b} in order to keep only m_b as in World_5.

Therefore, by transitivity, $\mathsf{World}_0 \approx \mathsf{World}_6$, i.e. the real and ideal worlds are indistinguishable $\mathsf{REAL}^{\sigma}_{\Pi,\hat{\mathbf{B}},\mathbf{Z}} \approx \mathsf{IDEAL}^{\sigma,\mathcal{F}_{OT}}_{\tilde{\Pi},\mathsf{Sim}_{\hat{\mathbf{B}}},\mathbf{Z}}$ which concludes this part of the proof.

Case 3: malicious receiver (Anboto). We consider now the case where the adversary corrupts the receiver Anboto, that is $\mathcal{A} = \hat{\mathbf{A}}$.

As in the previous case, without loss of generality we can decompose $\hat{\mathbf{A}}$ into three parts: $\hat{\mathbf{A}}_0$ will be the part running the Π_h protocol, $\hat{\mathbf{A}}_1$ will be the part running the (NI)ZK protocol, forwarding its internal state to $\hat{\mathbf{A}}_2$ that will be in charge of the rest of the protocol as pictured in Protocol 15. In the following, we define, for any bit b, $\mathcal{X}^{(b)}$ as the register containing the state $|\psi^{(b)}\rangle$, $\mathcal{W}^{(1-b)}$ as the sub-register of $\mathcal{X}^{(b)}$ containing the witness ω (so all but the first qubit) and $\mathcal{W}^{(1-b)}[1]$ as the first qubit in the register $\mathcal{W}^{(1-b)}$.

We formalize the above by defining the following hybrid worlds (see Protocols 15 to 19 for algorithmic descriptions):

i. $\mathsf{World}_0 := \mathsf{REAL}^{\sigma}_{\Pi,\hat{\mathbf{A}},\mathbf{Z}}$ is the real world (where Berain runs internally the (NI)ZK protocol $\mathbf{B}_{zk}$ with $\hat{\mathbf{A}}$).

ii. World_1 is like World_0 except that we replace Π_h with the ideal world ($\hat{\mathbf{A}}_0$

is replaced with $\mathsf{Sim}_{\hat{\mathbf{A}}_0}$, and $\mathbf{B}_h$ is now replaced with the ideal resource $\mathcal{F}_h^{\mathsf{Gen}}$ and the idealized party $\mathbf{B}_{h,I}$ sampling honestly $h \leftarrow \mathsf{Gen}(1^\lambda)$).

iii. World_2 is like World_1 except that we replace $\hat{\mathbf{A}}_1$ with the simulator of the (NI)ZK protocol interacting with $\mathcal{F}_{zk}^{\mathcal{R}}$, and integrate the ideal resource checking if there exists ω^{1-b} starting with a 1 in $\mathbf{B}$.

iv. World_3 is like World_2 except that it does not perform the $Z^{m_{1-b}}$ rotation.

v. For World_4, we can realize that in World_3, the code does not depend on m_{1-b} anymore. So we can reorganise the elements of World_3 to let the functionality provide m_b and discard m_{1-b}: we split $\mathbf{B}$ into three parts: the functionality $\mathcal{F}_{OT}$, a dummy Berain that just forwards m_0 and m_1 to $\mathcal{F}_{OT}$, and a simulator $\mathsf{Sim}_{\hat{\mathbf{A}}}$ that runs the code of $\mathbf{B}$ in the previous World (extracting b in the same way), except that it sends b to $\mathcal{F}_{OT}$ to get m_b. We also let $\mathsf{Sim}_{\hat{\mathbf{A}}}$ absorb $\hat{\mathbf{A}}$, $\mathbf{B}_{h,I}$,$\mathsf{Sim}_{\hat{\mathbf{A}}_0}$ and $\mathsf{Sim}_{\hat{\mathbf{A}}_1}$ (appropriately forwarding their input/outputs): This way $\mathsf{World}_4 = \mathsf{IDEAL}^{\sigma,\mathcal{F}_{OT}}_{\tilde{\Pi},\mathsf{Sim}_{\hat{\mathbf{A}}},\mathbf{Z}}$.

Protocol 15 Case 3 (malicious Anboto): World_0.

	$\hat{\mathbf{A}}(\sigma_{\hat{\mathbf{A}}})$		$\mathbf{B}(m_0, m_1)$	$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	$\hat{\mathbf{A}}_0$	$\longleftrightarrow$	Run $\mathbf{B}_h$ to obtain h	
2:	$\hat{\mathbf{A}}_1$	$\longleftrightarrow$	Run $\mathbf{B}_{zk}$ to obtain $\forall c, d, h_d^{(c)}$	
3:	$\hat{\mathbf{A}}_2$	$\xrightarrow{\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle}$	Do the measurements and rotations of the real protocol	
4:	$\hat{\mathbf{A}}_3$	$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$		
5:		$\xrightarrow{\mathsf{state}}$		**return** anything

First, we see that $\mathsf{World}_0 \approx \mathsf{World}_1$ because we assumed that the Π_h protocol $\mathsf{C-QSA}$ realizes $\mathcal{F}_h^{\mathsf{Gen}}$: If it were not the case, then we could easily break the $\mathsf{C-QSA}$ property of Π_h by merging basically all the procedure after the Π_h protocol inside $\mathbf{Z}$ to produce a new $\mathbf{Z}'$ able to attack $\mathcal{F}_h^{\mathsf{Gen}}$ with exactly the same probability. By practically the same argument, $\mathsf{World}_1 \approx \mathsf{World}_2$, because we assumed that the underlying protocol $\mathsf{C-QSA}$ realizes $\mathcal{F}_{zk}^{\mathcal{R}}$.

Protocol 16 Case 3 (malicious Anboto): World_1.

	$\hat{\mathbf{A}}(\sigma_{\hat{\mathbf{A}}})$		$\mathbf{B}(m_0, m_1)$	$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	$\mathsf{Sim}_{h,\hat{\mathbf{A}}_0}$	$\xleftarrow{h}$	$h \leftarrow \mathsf{Gen}(1^\lambda)$	
2:	$\hat{\mathbf{A}}_1$	$\longleftrightarrow$	Run $\mathbf{B}_{zk}$ to obtain $\forall c, d, h_d^{(c)}$	
3:	$\hat{\mathbf{A}}_2$	$\xrightarrow{\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle}$	Do the measurements and rotations of the real protocol	
4:	$\hat{\mathbf{A}}_3$	$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$		
5:		$\xrightarrow{\mathsf{state}}$		**return** anything

Protocol 17 Case 3 (malicious Anboto): World_2.

	$\hat{\mathbf{A}}(\sigma_{\hat{\mathbf{A}}})$		$\mathbf{B}(m_0, m_1)$	$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	$\mathsf{Sim}_{h,\hat{\mathbf{A}}_0}$	$\xleftarrow{h}$	$h \leftarrow \mathsf{Gen}(1^\lambda)$	
2:	$\mathsf{Sim}_{zk,\hat{\mathbf{A}}_1}$	$\xrightarrow{\forall c, d, \omega_d^{(c)}, h_d^{(c)}}$	Run $\mathcal{F}_{zk}^{\mathcal{R}}$	
3:	$\hat{\mathbf{A}}_2$	$\xrightarrow{\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle}$	Do the measurements and rotations of the real protocol	
4:	$\hat{\mathbf{A}}_3$	$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$		
5:		$\xrightarrow{\mathsf{state}}$		**return** anything

Protocol 18 Case 3 (malicious Anboto): World_3.

	$\hat{\mathbf{A}}(\sigma_{\hat{\mathbf{A}}})$		$\mathbf{B}(m_0, m_1)$	$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	$\mathsf{Sim}_{h,\hat{\mathbf{A}}_0}$	$\xleftarrow{h}$	$h \leftarrow \mathsf{Gen}(1^\lambda)$	
2:	$\mathsf{Sim}_{zk,\hat{\mathbf{A}}_1}$	$\xrightarrow{\forall c, d, \omega_d^{(c)}, h_d^{(c)}}$	Run $\mathcal{F}_{zk}^{\mathcal{R}}$	
3:	$\hat{\mathbf{A}}_2$	$\xrightarrow{\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle}$	Do the measurements and rotations of the real protocol, except $Z^{m_{1-b}}$	
4:	$\hat{\mathbf{A}}_3$	$\xleftarrow{\forall c, w^{(c)}, z^{(c)}}$		
5:		$\xrightarrow{\mathsf{state}}$		**return** anything

Protocol 19 Case 3 (malicious Anboto): World_4.

	$\mathsf{Sim}_{\hat{\mathbf{A}}(\sigma_{\hat{\mathbf{A}}})}$		$\tilde{\mathbf{B}}(m_0, m_1)$ and $\mathcal{F}_{OT}$	$\mathbf{Z}(\sigma_{\mathbf{Z}})$
1:	$h \leftarrow \mathsf{Gen}(1^\lambda)$			
2:	$\mathsf{state}_0 \leftarrow \mathsf{Sim}_{h,\hat{\mathbf{A}}_0}(\sigma_{\hat{\mathbf{A}}}, h)$			
3:	$\forall c(d, \omega_d^{(c)}, h_d^{(c)}, \mathsf{state}_1)$ $\leftarrow \mathsf{Sim}_{zk,\hat{\mathbf{A}}_1}(\mathsf{state}_0)$			
4:	Run $\mathcal{F}_{zk}^{\mathcal{R}}$	$\xrightarrow{b}$		
5:	$(\lvert\psi^{(0)}\rangle, \lvert\psi^{(1)}\rangle, \mathsf{state}_2)$ $\leftarrow \hat{\mathbf{A}}_2(\mathsf{state}_1)$	$\xleftarrow{m_b}$		
6:	Do the measurements and rotations of the real protocol, except $Z^{m_{1-b}}$			
7:	$\mathsf{state}_3 \leftarrow \hat{\mathbf{A}}_3(\forall c, \omega^{(c)}, z^{(c)}, \mathsf{state}_2)$	$\xrightarrow{\mathsf{state}_3}$		**return** anything

Proving that $\mathsf{World}_2 \approx \mathsf{World}_3$ is more technical, but intuitively the goal is to show that the state sent by Anboto is close in trace distance to a state in the computational basis (otherwise we can break the collision resistance property of h), and therefore a Z rotation does not significantly disturb the state. For now we leave it as a claim that $\mathsf{World}_2 \approx \mathsf{World}_3$ whenever h is injective or computationally collision resistant, and leave the proof for later.

Claim 5.3.1. *When h is injective, then* $\mathsf{World}_2 \approx \mathsf{World}_3$*.*

Claim 5.3.2. *When h is collision resistant, then* $\mathsf{World}_2 \approx \mathsf{World}_3$*.*

Finally, it is easy to see that $\mathsf{World}_3 = \mathsf{World}_4$ since they are actually exactly the same quantum map, except that we attribute the operations to different parties. By defining $\mathsf{Sim}_{\hat{\mathbf{A}},\mathbf{Z}}$ as the block composed of all elements on the left of the ideal resource, we have $\mathsf{World}_5 = \mathsf{IDEAL}^{\sigma,\mathcal{F}_{OT}}_{\tilde{\Pi},\mathsf{Sim}_{\hat{\mathbf{A}}},\mathbf{Z}}$.

By transitivity, we have $\mathsf{World}_0 \approx \mathsf{World}_4$, i.e. the real and ideal worlds are indistinguishable $\mathsf{REAL}^{\sigma}_{\Pi,\hat{\mathbf{A}},\mathbf{Z}} \approx \mathsf{IDEAL}^{\sigma,\mathcal{F}_{OT}}_{\tilde{\Pi},\mathsf{Sim}_{\hat{\mathbf{A}}},\mathbf{Z}}$, which concludes the proof. □

In order to prove the claims, we need to formalize the intuition that the state sent by Alice is close in trace distance to a state in the computational basis. We will first associate a quantity β to each run, show that this β is linked to the probability of finding a collision and to the trace distance to the measured state, and finally we show that the average value of β must be negligible, like the average trace distance between the two worlds. To that aim, it is handy to consider, for a fixed run, a (normalized) purification $|\psi\rangle_{\mathcal{T},\mathbf{B},E} = |t\rangle \sum_{x,y} \beta^{(1-b)}_{x,y} |x\rangle|y\rangle$ of the states $|\psi^{(0)}\rangle$ and $|\psi^{(1)}\rangle$ sent by $\hat{\mathbf{A}}$ and partially measured[3] by $\mathbf{B}$ with the result of the tests it the register $\mathcal{T}$, including any potential entanglement with the adversary or environment by adding a third register $\mathcal{E}$ (we also put in this register the internal memory of $\hat{\mathbf{A}}$ right after she sent the state to $\mathbf{B}$). Moreover, we can assume without loss of generality that $\hat{\mathbf{A}}$ does nothing before receiving the measurement outcomes sent by Berain, by simply postponing in time its actions.

Proof of Claim 5.3.1. First, if $t = 0$ (invalid test), then the remaining actions in World_2 and World_3 are identical. Now, if $t = 1$, because h is injective there

[3] We are referring to the operations where Berain extracts the witness, where it checks that the quantum register containing the witnesses in superposition starts with a 0 and that they are valid pre-images. The result of these tests ($t = 1$ if and only if all tests pass) is put in a new register $\mathcal{T}$.

exists at most one $x^{(1-b)}$ such $x^{(1-b)}$ has a 0 at the second position and $h(x^{(1-b)}) \in \{h_0^{(1-b)}, h_1^{(1-b)}\}$ (we already extracted above one pre-image of $h_0^{(1-b)}$ or $h_1^{(1-b)}$ with a 1 at the second position in the ZK protocol, and by injectivity of h there are at most two pre-images to this set, so a single image can have a 0 at the second position). Because Berain measured exactly that the first register is a pre-image of $\{h_0^{(1-b)}, h_1^{(1-b)}\}$ (since $t = 1$), the state $|\psi\rangle_{\mathbf{B},E}$ will collapse into $|x^{(1-b)}\rangle(\sum_y \beta_{x,y}^{(1-b)}|y\rangle)$ (up to a renormalisation factor). Therefore, since $Z^m|x\rangle = |x\rangle$, applying $Z_2^{m_{1-b}}$ or not does not change the state at all (this is true for any run), leading to $\mathsf{World}_2 = \mathsf{World}_3$. □

Proof of Claim 5.3.1. During a valid extraction of the witness, we found an element $x_{1-s}^{(1-b)}$ whose witness part starts with a 1 ($x_{1-s}^{(1-b)}[2] = 1$) such that $h(x_{1-s}^{(1-b)}) \in \{h_0^{(1-b)}, h_1^{(1-b)}\}$: Let $x^* := x_s^{(1-b)}$ be the other part of the witness such that $h(x^*) \in \{h_0^{(1-b)}, h_1^{(1-b)}\}$ (if both of them start with a 1, we can choose arbitrarily). Then, there exists $\beta \in [0, 1]$ (possibly equal to 1 if x^* starts with a 1), a normalized pure state $|\phi^*\rangle$, and a normalized pure state $|\phi\rangle$ such that $\mathrm{Tr}((|x^*\rangle\langle x^*| \otimes I)|\phi\rangle\langle\phi|) = 0$ such that:

$$|\psi\rangle_{\mathcal{T},\mathbf{B},E} = |t\rangle(\sqrt{1-\beta}|x^*\rangle|\phi^*\rangle + \sqrt{\beta}|\phi\rangle)\,. \tag{5.17}$$

We observe now that the probability of finding a collision is greater than $t\beta$. First, if $t = 0$ (the test fails), then $t\beta = 0$ so this is obviously true. Now, if $t = 1$, this can be seen by first remarking that if we measure the register $\mathbf{B}$ of $|\psi\rangle_{\mathbf{B},E}$, we get an outcome x: since $\mathrm{Tr}((|x^*\rangle\langle x^*| \otimes I)|\phi\rangle\langle\phi|) = 0$, x equals x^* with probability $\sqrt{1-\beta}^2 = 1 - \beta$, i.e. $x \neq x^*$ with probability β. Moreover, because $t = 1$, we know that $x[2] = 0$ and $h(x) \in \{h_0^{(1-b)}, h_1^{(1-b)}\}$. In the first case, if $h(x) = h_{1-s}^{(1-b)}$, then $x \neq x_{1-s}^{(1-b)}$ since $x[2] = 1$ and $x_{1-s}^{(1-b)}[2] = 0$ so $(x, x_{1-s}^{(1-b)})$ is a collision (reminder: $x_{1-s}^{(1-b)}$ was extracted by the simulator during the ZK protocol). In the second case, if $h(x) = h_s^{(1-b)}$, because with probability β we have $x \neq x^*$, (x, x^*) forms a collision with probability β. Therefore we can find a collision with probability greater than $\beta = t\beta$.

Moreover, we observe that the trace distance between the states $|\psi\rangle_{\mathcal{T},\mathbf{B},E}$ and $Z_{\mathbf{B},2}^{tm_{1-b}}|\psi\rangle_{\mathbf{B},E}$ is smaller than $2t\sqrt{\beta}$. First, if $t = 0$, both states are strictly equal, so their trace distance is 0. If $t = 1$

$$\mathrm{TD}(|\psi\rangle_{\mathcal{T},\mathbf{B},E}, (Z_{\mathbf{B},2}^{tm_{1-b}}|\psi\rangle_{\mathcal{T},\mathbf{B},E})) = \mathrm{TD}(|t\rangle|\psi\rangle_{\mathbf{B},E}, |t\rangle Z_{\mathbf{B},2}^{tm_{1-b}}|\psi\rangle_{\mathbf{B},E}) \tag{5.18}$$

$$\begin{aligned}&\leq \mathrm{TD}(|\psi\rangle_{\mathbf{B},E}, |x^*\rangle|\phi^*\rangle) + \mathrm{TD}(|x^*\rangle|\phi^*\rangle, Z_2^{m_{1-b}}|x^*\rangle|\phi^*\rangle)\\ &\quad + \mathrm{TD}(Z_2^{m_{1-b}}|x^*\rangle|\phi^*\rangle, Z_2^{m_{1-b}}|\psi\rangle_{\mathbf{B},E})\end{aligned} \tag{5.19}$$

$$\leq \sqrt{\beta} + 0 + \mathrm{TD}(|x^*\rangle|\phi^*\rangle, |\psi\rangle_{\mathbf{B},E}) \tag{5.20}$$

$$\leq 2\sqrt{\beta} = 2t\sqrt{\beta}\,. \tag{5.21}$$

where we used the triangle inequality, the fact that on states in the computational basis, $Z|x\rangle = |x\rangle$, Lemma 5.2.10, and that the trace distance is preserved under unitary transforms.

We prove now that $\mathsf{World}_2 \approx_\alpha \mathsf{World}_3$, where the error parameter is $\alpha := \mathbb{E}_{2t\sqrt{\beta_\psi}}\left[|t\rangle|\psi\rangle_{\mathbf{B},E} \leftarrow \xi_0(\sigma)\right]$, and show later that α must be negligible. First, we remark that by stopping the worlds right before the Z rotations, we can define two binary POVM[4] ξ_0 (taking as input (a purification of) σ_λ and outputting the state $|\psi\rangle_{\mathcal{T},\mathbf{B},E} = |t\rangle|\psi\rangle_{\mathbf{B},E}$ defined above) and ξ_1 (performing the rest Z^{m_b} rotation, the measurement in the H basis, the adversary $\hat{\mathbf{A}}_2$ and $\mathbf{Z}$) such that World_3 is the sequential composition of ξ_0 and ξ_1, and World_2 is the sequential composition of ξ_0, $Z_2^{tm_{1-b}}$ and ξ_1. This way, we can write:

$$|\Pr(\mathsf{World}_3 = 1) - \Pr(\mathsf{World}_2 = 1)| \tag{5.22}$$

$$\begin{aligned}&= \Bigg| \Pr_{|t\rangle|\psi\rangle_{\mathbf{B},E} \leftarrow \xi_0(\sigma)}\Big(\xi_1(|t\rangle|\psi\rangle_{\mathbf{B},E}) = 1\Big)\\ &\quad - \Pr_{|t\rangle|\psi\rangle_{\mathbf{B},E} \leftarrow \xi_0(\sigma)}\Big(\xi_1(|t\rangle(Z_{\mathbf{B},2}^{tm_{1-b}}|\psi\rangle_{\mathbf{B},E}) = 1\Big)\Bigg|\end{aligned} \tag{5.23}$$

$$\begin{aligned}&= \Bigg| \int_{|t\rangle|\psi\rangle_{\mathbf{B},E}} \Pr\Big(\xi_0(\sigma) = |t\rangle|\psi\rangle_{\mathbf{B},E}\Big)\Big(\Pr\Big(\xi_1(|t\rangle|\psi\rangle_{\mathbf{B},E}) = 1\Big)\\ &\quad - \Pr\Big(\xi_1(|t\rangle(Z_{\mathbf{B},2}^{tm_{1-b}}|\psi\rangle_{\mathbf{B},E}))\Big)\Big)\Bigg|\end{aligned} \tag{5.24}$$

$$\begin{aligned}&\leq \int_{|t\rangle|\psi\rangle_{\mathbf{B},E}} \Pr\Big(\xi_0(\sigma) = |t\rangle|\psi\rangle_{\mathbf{B},E}\Big)\Big|\Pr\Big(\xi_1(|t\rangle|\psi\rangle_{\mathbf{B},E}) = 1\Big)\\ &\quad - \Pr\Big(\xi_1\Big(|t\rangle(Z_{\mathbf{B},2}^{tm_{1-b}}|\psi\rangle_{\mathbf{B},E})\Big) = 1\Big)\Big|\end{aligned} \tag{5.25}$$

[4] We slightly abuse notations, as technically a POVM is not a map but a set of projectors (one for each outcome), so we define $\Pr(\xi_0(\rho) = 1) = \mathrm{Tr}(\xi_0\rho)$.

$$= \mathop{\mathbb{E}}_{\xi_0(\sigma)=|t\rangle|\psi\rangle_{\mathbf{B},E}} \Big[\Big| \Pr\Big(\xi_1(|t\rangle|\psi\rangle_{\mathbf{B},E}) = 1\Big) \\ - \Pr\Big(\xi_1(|t\rangle(Z_{\mathbf{B},2}^{tm_{1-b}}|\psi\rangle_{\mathbf{B},E})) = 1\Big)\Big|\Big] . \tag{5.26}$$

Moreover, by the operational interpretation of the trace distance, since ξ_1 is a POVM, we have that

$$\Big|\Pr\Big(\xi_1(|t\rangle|\psi\rangle_{\mathbf{B},E}) = 1\Big) - \Pr\Big(\xi_1(|t\rangle(Z_{\mathbf{B},2}^{tm_{1-b}}|\psi\rangle_{\mathbf{B},E})) = 1\Big)\Big| \tag{5.27}$$

$$\leq \mathrm{TD}(|\psi\rangle_{\mathcal{T},\mathbf{B},E}, Z_{\mathbf{B},2}^{tm_{1-b}}|\psi\rangle_{\mathcal{T},\mathbf{B},E}) \tag{5.28}$$

$$\leq 2t\sqrt{\beta_\psi} \, , \tag{5.29}$$

where in the last line we used eq. (5.21). Note that β is different for any value of $|\psi\rangle_{\mathcal{T},\mathbf{B},E}$, hence the notation β_ψ. By injecting the above equation into eq. (5.26), we get

$$\Big|\Pr(\mathsf{World}_3 = 1) - \Pr(\mathsf{World}_2 = 1)\Big| \leq \mathop{\mathbb{E}}_{|t\rangle|\psi\rangle_{\mathbf{B},E} \leftarrow \xi_0(\sigma)} [2t\sqrt{\beta_\psi}] = \alpha \, . \tag{5.30}$$

Finally, we prove that α is negligible by reducing it to the probability of finding a collision. We have already shown above that given $|t\rangle|\psi\rangle_{\mathbf{B},E}$, there is a procedure P_h to find a collision with probability greater than $t\beta_\psi$ (but the initial state might depend on h). From that, we can define the algorithm that first runs $\xi_0(\sigma_\lambda)$ (σ_λ is now independent of h, and h is sampled according to $\mathsf{Gen}(1^\lambda)$ in ξ_0 as expected), then P_h. The probability of success for this procedure is therefore:

$$\alpha' := \mathop{\mathbb{E}}_{|t\rangle|\psi\rangle_{\mathbf{B},E} \leftarrow \xi_0(\sigma)} [t\beta_\psi] \, . \tag{5.31}$$

Since h is collision resistant, we have $\alpha' = \mathsf{negl}(\lambda)$. Moreover, by defining $f(x) = \sqrt{2x}$, f is concave, and therefore using Jensen's inequality we get:

$$\alpha = \mathop{\mathbb{E}}_{|t\rangle|\psi\rangle_{\mathbf{B},E} \leftarrow \xi_0(\sigma)} [2t\sqrt{\beta_\psi}] \tag{5.32}$$

$$\leq 2 \mathop{\mathbb{E}}_{|t\rangle|\psi\rangle_{\mathbf{B},E} \leftarrow \xi_0(\sigma)} [\sqrt{t\beta_\psi}] \tag{5.33}$$

$$\leq 2 \sqrt{\mathop{\mathbb{E}}_{|t\rangle|\psi\rangle_{\mathbf{B},E} \leftarrow \xi_0(\sigma)} [t\beta_\psi]} \tag{5.34}$$

$$= 2\sqrt{\alpha'} \tag{5.35}$$

$$= \mathsf{negl}(\lambda) \, . \tag{5.36}$$

□

5.4 Composability of Unruh's NIZK protocol

In this section we will prove that Unruh's [Unr15] online extractable non-interactive zero-knowledge (NIZK) protocol quantum stand-alone $\mathsf{C-QSA}$) realizes the ideal zero-knowledge functionality $\mathcal{F}_{zk}^{\mathcal{R}}$, in the quantum random oracle model (QROM).

Note that the polynomial-time QIM prover P and verifier V from [Unr15, Fig. 1] have access to two random oracles, G and H, which can be queried in superposition by both parties (for simplicity we will just refer to a single oracle H). We will denote Unruh's polynomial-time two-party NIZK protocol by $\Pi_{zk}^{H} = (\mathsf{P}, \mathsf{V})$, to stress the interaction between two machines and the trusted random oracles $\mathsf{P} \overset{H}{\leftrightsquigarrow} \mathsf{V}$.

In the original article, the protocol Π_{zk}^{H} is proven to be complete, zero-knowledge and (even simulation-sound) online-extractable. We recall here these definitions for clarity (note that we assume they also hold against non-uniform adversaries).

Definition 5.4.1 ([Unr15, Definitions 1, 2 and 3])**.** *Let* P *and* V *be two polynomial-time QIMs and let* $\Pi_{zk}^{H} = (\mathsf{P}, \mathsf{V})$ *be a non-interactive proof system. Let* $\mathcal{R}$ *be a efficiently decidable relation defining the language of our proof system.*

i. ***Completeness:*** Π_{zk}^{H} *is* complete *if and only if for any* QPT *oracle algorithm* $\mathcal{A}$ *and advice* $\{\sigma_\lambda\}_{\lambda\in\mathbb{N}}$*,*

$$\Pr\Big((x,w)\in\mathcal{R}\wedge y=0\Big|H\leftarrow \mathsf{ROdist}, (x,w)\leftarrow\mathcal{A}^{H}(\sigma_\lambda), \\ y\leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(x,w)\overset{H}{\leftrightsquigarrow}\mathsf{V}\rangle\Big)\leq \mathsf{negl}(\lambda)\,. \quad (5.37)$$

ii. ***Zero-knowledge:*** Π_{zk}^{H} *is* zero-knowledge *if and only if there exists a polynomial-time simulator* $\mathsf{Sim} = (\mathsf{S}_{\text{init}}, \mathsf{S}_{\mathsf{P}})$ *such that for every* QPT *oracle algorithm* $\mathcal{A}$ *and advice* $\{\sigma_\lambda\}_{\lambda\in\mathbb{N}}$*,*

$$\Big|\Pr\Big(H\leftarrow \mathsf{ROdist}, z\leftarrow\mathcal{A}^{H,\mathsf{P}}(\sigma_\lambda)\Big|z=1\Big) \\ -\Pr\Big(H\leftarrow \mathsf{S}_{\text{init}}, z\leftarrow\mathcal{A}^{H,\mathsf{S}_{\mathsf{P}}}(\sigma_\lambda)\Big|z=1\Big)\Big|\leq \mathsf{negl}(\lambda)\,. \quad (5.38)$$

Since in the quantum setting we cannot allow the simulator to learn the input for each query, because this can be done in superposition, *here the simulator S_{init} outputs a circuit describing a classical function representing the initial random oracle instead. We assume that both S_{init} and S_{P} have access to the polynomial upper bound on the runtime of $\mathcal{A}$.*

iii. ***Online-extractability:*** *Π_{zk}^{H} is* online extractable *with respect to S_{init} if and only if there exists a polynomial-time extractor E such that for any QPT oracle algorithm $\mathcal{A}$ and advice $\{\sigma_\lambda\}_{\lambda\in\mathbb{N}}$,*

$$\Pr\Big(y = 1 \wedge (x,w) \notin \mathcal{R} \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x,\pi) \leftarrow \mathcal{A}^{H}(\sigma_\lambda), y \leftarrow \mathsf{V}^{H}(x,\pi), w \leftarrow \mathsf{E}(H,x,\pi)\Big) \leq \mathsf{negl}(\lambda)\,. \quad (5.39)$$

We assume that both S_{init} and E have access to the polynomial upper bound on the runtime of $\mathcal{A}$.

Note that Unruh's original definition only considers uniform adversaries, whilst here we assume that the above conditions hold for the protocol when the adversary receives advice.

We can now prove the main theorem of this section.

Theorem 5.4.2. *Let H be a quantum random oracle. The non-interactive protocol $\Pi_{zk}^{H} = (\mathsf{P},\mathsf{V})$ from [Unr15] quantum stand-alone realizes the classical zero-knowledge functionality $\mathcal{F}_{zk}^{\mathcal{R}}$.*

Proof. We define trivially the dummy parties $\widetilde{\Pi}_{zk} = (\widetilde{\mathsf{P}},\widetilde{\mathsf{V}})$ that forward the inputs/outputs to/from $\mathcal{F}_{zk}^{\mathcal{R}}$. We want to show that for any (poly-)time adversary $\mathcal{A}$ there exists a (poly-)time simulator Sim such that, for any poly-time distinguisher $\mathbf{Z}$ and advice σ, we have $\mathsf{REAL}^{\sigma}_{\Pi_{zk},\mathcal{A},\mathbf{Z}} \approx \mathsf{IDEAL}^{\sigma,\mathcal{F}_{zk}^{\mathcal{R}}}_{\widetilde{\Pi}_{zk},\mathsf{Sim}_{\mathcal{A}},\mathbf{Z}}$. We will split the proof depending on the parties that the static adversary $\mathcal{A}$ corrupts (nobody, the prover P or the verifier V).

Note that in the non-interactive protocol Π_{zk}^{H} the prover rejects if it receives a non-valid witness $(x,w) \notin \mathcal{R}$, which is essential.

Case 1: correctness (no corrupted party). Let $\lambda \in \mathbb{N}$. For any bipartite input state $\sigma_\lambda^{\mathsf{P},\mathbf{Z}} \in \mathrm{D}(\mathcal{S}_\lambda \otimes \mathsf{R}_\lambda)$, with R_λ an arbitrary reference system, we want to show that for any environment $\mathbf{Z}$ the probability of distinguishing

Π^H_{zk} from the ideal functionality $\widetilde{\Pi}_{zk}$ is negligible:

$$\Big|\Pr\Big(\mathbf{Z}(y,\sigma^{\mathbf{Z}}_\lambda)=1\Big|H\leftarrow \mathsf{ROdist}, y\leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(\sigma^{\mathsf{P}}_\lambda)\overset{H}{\leftrightsquigarrow}\mathsf{V}\rangle\Big) \\ -\Pr\Big(\mathbf{Z}(y,\sigma^{\mathbf{Z}}_\lambda)=1\Big|H\leftarrow \mathsf{ROdist}, y\leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \widetilde{\mathsf{P}}(\sigma^{\mathsf{P}}_\lambda)\overset{\mathcal{F}}{\leftrightsquigarrow}\widetilde{\mathsf{V}}\rangle\Big)\Big|\leq \mathsf{negl}(\lambda)\,. \tag{5.40}$$

In order to prove the above inequality we are interested in rewriting the probability of the distinguisher outputting 1 in terms of the input/output of the interaction. Since the prover expects a classical message, we can model $\sigma^{\mathsf{P},\mathbf{Z}}_\lambda$ as a quantum instrument:

$$\sigma^{\mathsf{P},\mathbf{Z}}_\lambda=\sum_{x,w}p_{x,w}|x,w\rangle\langle x,w|\otimes\sigma_{\lambda,x,w}\,. \tag{5.41}$$

This allows to write the LHS of eq. (5.40) as

$$\sum_{x,w}p_{x,w}\Pr\Big(\mathbf{Z}(y,\sigma_{\lambda,x,w})=1\Big|H\leftarrow \mathsf{ROdist}, y\leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(x,w)\overset{H}{\leftrightsquigarrow}\mathsf{V}\rangle\Big) \\ =\sum_{b\in\{0,1\}}\sum_{x,w}\Big(p_{x,w}\Pr(\mathbf{Z}(b,\sigma_{\lambda,x,w})=1) \\ \cdot\Pr\Big(y=b\Big|H\leftarrow \mathsf{ROdist}, y\leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(x,w)\overset{H}{\leftrightsquigarrow}\mathsf{V}\rangle\Big)\Big)\,. \tag{5.42}$$

For the dummy protocol $\widetilde{\Pi}_{zk}$ the above equation has a very simple form since if $(x,w)\in\mathcal{R}$ (resp. $(x,w)\notin\mathcal{R}$), then $\widetilde{\Pi}_{zk}$ will output 1 (resp. 0) with probability 1. Therefore, for any input state $\sigma^{\mathsf{P},\mathbf{Z}}_\lambda$ we can simplify the RHS of eq. (5.40) to

$$\sum_{(x,w)\in\mathcal{R}}p_{x,w}\Pr(\mathbf{Z}(1,\sigma_{\lambda,x,w})=1)+\sum_{(x,w)\notin\mathcal{R}}p_{x,w}\Pr(\mathbf{Z}(0,\sigma_{\lambda,x,w})=1)\,. \tag{5.43}$$

For the honest protocol Π^H_{zk}, note that for invalid witnesses $(x,w)\notin\mathcal{R}$ the honest prover P^H will also always reject, therefore for $(x,w)\notin\mathcal{R}$:

$$\Pr\Big(y=1\Big|H\leftarrow \mathsf{ROdist}, y\leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(x,w)\overset{H}{\leftrightsquigarrow}\mathsf{V}\rangle\Big)=0\,, \tag{5.44}$$

thus for an arbitrary mixture of invalid witnesses $\sum_{(x,w)\notin\mathcal{R}}q_{x,w}|x,w\rangle\langle x,w|$, we will have that

$$\sum_{(x,w)\notin\mathcal{R}}q_{x,w}\Pr\Big(y=1\Big|H\leftarrow \mathsf{ROdist}, y\leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(x,w)\overset{H}{\leftrightsquigarrow}\mathsf{V}\rangle\Big)=0\,, \tag{5.45}$$

and consequently

$$\sum_{(x,w)\notin\mathcal{R}} q_{x,w} \Pr\Big(y=0 \Big| H \leftarrow \mathsf{ROdist}, y \leftarrow \mathsf{OUT_V}\langle \mathsf{P}(x,w) \overset{H}{\leftrightsquigarrow} \mathsf{V}\rangle\Big) = 1\,. \tag{5.46}$$

The case of valid witnesses is also easy as we know from completeness eq. (5.37) that given any input state $(x,w)\in\mathcal{R}$, if we pick the constant algorithm $\mathcal{A}\colon \sigma_\lambda^{\mathsf{P}} \mapsto (x,w)$, then

$$\Pr\Big(y=0 \Big| H \leftarrow \mathsf{ROdist}, y \leftarrow \mathsf{OUT_V}\langle \mathsf{P}(x,w) \overset{H}{\leftrightsquigarrow} \mathsf{V}\rangle\Big) \leq \mathsf{negl}(\lambda)\,, \tag{5.47}$$

thus for an arbitrary mixture of valid witnesses $\sum_{(x,w)\in\mathcal{R}} q_{x,w}|x,w\rangle\langle x,w|$, we will have that

$$\sum_{(x,w)\in\mathcal{R}} q_{x,w} \Pr\Big(y=0 \Big| H \leftarrow \mathsf{ROdist}, y \leftarrow \mathsf{OUT_V}\langle \mathsf{P}(x,w) \overset{H}{\leftrightsquigarrow} \mathsf{V}\rangle\Big) \leq \mathsf{negl}(\lambda)\,, \tag{5.48}$$

and consequently

$$\sum_{(x,w)\in\mathcal{R}} q_{x,w} \Pr\Big(y=1 \Big| H \leftarrow \mathsf{ROdist}, y \leftarrow \mathsf{OUT_V}\langle \mathsf{P}(x,w) \overset{H}{\leftrightsquigarrow} \mathsf{V}\rangle\Big) \geq 1-\mathsf{negl}(\lambda)\,. \tag{5.49}$$

We can combine the above eqs. (5.45), (5.46), (5.48) and (5.49) to obtain the desired inequality eq. (5.40) for any input $\sigma_\lambda^{\mathsf{P},\mathbf{Z}}$ and any distinguisher $\mathbf{Z}$ by noting that for any received input $b\in\{0,1\}$ and advice $\sigma_{\lambda,x,w}$, the probability of the distinguisher outputing 1 is bounded

$$\Pr(\mathbf{Z}(b,\sigma_{\lambda,x,w})=1) \leq 1\,, \tag{5.50}$$

and therefore by developing eq. (5.40) in terms of the advice as in eq. (5.42), we can bound the difference by

$$
\begin{aligned}
\leq & \left| \sum_{(x,w)\in\mathcal{R}} p_{x,w} \Pr\Big(y=0 \Big| H \leftarrow \mathsf{ROdist}, y \leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(x,w) \overset{H}{\leftrightsquigarrow} \mathsf{V}\rangle\Big) \right| \\
& + \left| \sum_{(x,w)\in\mathcal{R}} p_{x,w} \Big(\Pr\Big(y=1 \Big| H \leftarrow \mathsf{ROdist}, y \leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(x,w) \overset{H}{\leftrightsquigarrow} \mathsf{V}\rangle\Big) - 1\Big) \right| \\
& + \left| \sum_{(x,w)\notin\mathcal{R}} p_{x,w} \Big(\Pr\Big(y=0 \Big| H \leftarrow \mathsf{ROdist}, y \leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(x,w) \overset{H}{\leftrightsquigarrow} \mathsf{V}\rangle\Big) - 1\Big) \right| \\
& + \left| \sum_{(x,w)\notin\mathcal{R}} p_{x,w} \Pr\Big(y=1 \Big| H \leftarrow \mathsf{ROdist}, y \leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{P}(x,w) \overset{H}{\leftrightsquigarrow} \mathsf{V}\rangle\Big) \right| \\
\leq & \mathsf{negl}(\lambda)\,.
\end{aligned}
\tag{5.51}
$$

Case 2: malicious prover (Anboto). If the adversary corrupts the prover $\mathcal{A} = \hat{\mathsf{P}}$, we will use online-extractability to construct the desired simulator. Let $\mathsf{Sim}_{\hat{\mathsf{P}}}$ be the simulator from Protocol 20, we will now prove that no distinguisher can differentiate between the real protocol with corrupted prover Anboto and the ideal functionality with simulator $\mathsf{Sim}_{\hat{\mathsf{P}}}$. This proof relies on the closeness of the following hybrid worlds (see figures 5.5 to 5.9 for a graphical depiction):

Protocol 20 Simulator $\mathsf{Sim}_{\hat{\mathsf{P}}} := (\mathsf{S}, \mathsf{E}, \mathsf{S}_{\text{init}})$ for malicious prover $\mathcal{A} = \hat{\mathsf{P}}$.

```
1: $\mathsf{S}$ initializes $\hat{\mathsf{P}}$ with whatever input state it receives.
2: $\mathsf{S}$ obtains $(x, \pi)$ from $\hat{\mathsf{P}}$.
3: $\mathsf{S}$ initializes $\mathsf{E}$ with $(x, \pi)$ and the description of the oracle $H$ given by
   the simulator $\mathsf{S}_{\text{init}}$.
4: $\mathsf{E}(x, \pi, H)$ extracts a witness $w$ or an abort message $\bot$, sends it to $\mathsf{S}$.
5: $\mathsf{S}$ sends $(x, w)$ or $\bot$ to $\mathcal{F}_{zk}^{\mathcal{R}}$.
```

i. $\mathsf{World}_0 := \mathsf{IDEAL}^{\sigma, \mathcal{F}_{zk}^{\mathcal{R}}}_{\widetilde{\Pi}_{zk}, \mathsf{Sim}_{\hat{\mathsf{P}}}, \mathbf{Z}}$ is the ideal world. Consists of one output by the functionality (forwarded by the dummy verifier) which is accepting if the witness obtained by the extractor in line 4 is valid, i.e., $(x, w) \in \mathcal{R}$.

ii. World_1 is like World_0 except that we substitute the dummy verifier $\widetilde{\mathsf{V}}$ by a merge of Unruh's verifier V and the simulator. In particular, we replace $\widetilde{\mathsf{V}}$ by a verifier V_1 that forwards the proof from the simulator S to the extractor E. If the extractor provides a witness w, then V_1 accepts and outputs $y := x$, else aborts and outputs $y := \bot$.

iii. World_2 is like World_1 except that we drop the extractor E (as it is only being used to check the proof) and the dummy verifier V_1 which is only forwarding information, and we use the verifier V to perform the check of the proof received by the simulator S instead.

iv. World_3 is like World_2 except that we drop the simulator S as it is only forwarding the information to the verifier.

v. World_4 differs from World_3 in that we replace the simulator S_{init} by the oracle H that is simulating. Note that now $\mathsf{World}_4 := \mathsf{REAL}^{\sigma}_{\Pi_{zk},\hat{\mathsf{P}},\mathbf{Z}}$.

The similarity of the first two worlds, $\mathsf{World}_0 \approx \mathsf{World}_1$, is a consequence of online-extractability. More precisely, if we could distinguish these two worlds, there would exit an input $\sigma^{\mathsf{P}}_\lambda$ such that

$$\Big|\Pr\Big(\mathbf{Z}(y,\sigma^{\mathbf{Z}}_\lambda)=1\Big|H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\widetilde{\mathsf{V}}}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma^{\mathsf{P}}_\lambda) \overset{\mathcal{F}}{\leftrightsquigarrow} \widetilde{\mathsf{V}}\rangle\Big)$$
$$-\Pr\Big(\mathbf{Z}(y,\sigma^{\mathbf{Z}}_\lambda)=1\Big|H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{V_1}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma^{\mathsf{P}}_\lambda) \leftrightsquigarrow \mathsf{V}_1\rangle\Big)\Big| > \mathsf{negl}(\lambda)\,. \tag{5.52}$$

In order to work with them jointly we expand them in terms of the verifier accepting/rejecting $y \in \{x,\bot\}$. Let us denote by $\delta(\mathsf{V})$ the function outputting 1 if the verifier V accepts, and 0 otherwise.

$$\Big|\Pr\Big(\mathbf{Z}(y,\sigma^{\mathbf{Z}}_\lambda)=1\Big|H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\widetilde{\mathsf{V}}}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma^{\mathsf{P}}_\lambda) \overset{\mathcal{F}}{\leftrightsquigarrow} \widetilde{\mathsf{V}}\rangle\Big)$$
$$-\Pr\Big(\mathbf{Z}(y,\sigma^{\mathbf{Z}}_\lambda)=1\Big|H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\mathsf{V}_1}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma^{\mathsf{P}}_\lambda) \leftrightsquigarrow \mathsf{V}_1\rangle\Big)\Big| \tag{5.53}$$
$$=\Big|\sum_{b\in\{x,\bot\}} \Pr\big(\mathbf{Z}(b,\sigma^{\mathbf{Z}}_\lambda)=1\big)$$
$$\cdot\Pr\Big(y=b\Big|H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\widetilde{\mathsf{V}}}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma^{\mathsf{P}}_\lambda) \overset{\mathcal{F}}{\leftrightsquigarrow} \widetilde{\mathsf{V}}\rangle\Big)$$
$$-\sum_{b\in\{x,\bot\}} \Pr\big(\mathbf{Z}(b,\sigma^{\mathbf{Z}}_\lambda)=1\big) \tag{5.54}$$
$$\cdot\Pr\Big(y=b\Big|H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{V_1}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma^{\mathsf{P}}_\lambda) \leftrightsquigarrow \mathsf{V}_1\rangle\Big)\Big|$$

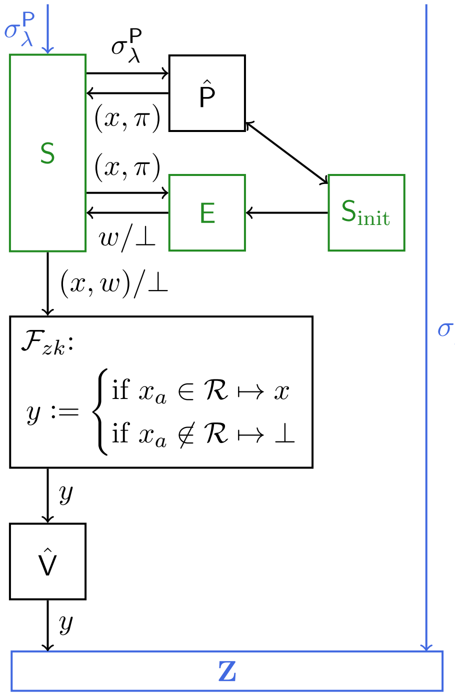

$\sigma_\lambda^{\mathsf{P}}$
$\sigma_\lambda^{\mathsf{P}}$
S
$\hat{\mathsf{P}}$
(x,π)
(x,π)
E
S_{init}
$w/\bot$
$(x,w)/\bot$
$\mathcal{F}_{zk}$: $y := \begin{cases}\text{if } x_a \in \mathcal{R} \mapsto x \\ \text{if } x_a \notin \mathcal{R} \mapsto \bot\end{cases}$
y
$\hat{\mathsf{V}}$
y
$\sigma_\lambda^{\mathbf{Z}}$
Z

Figure 5.5: World_0

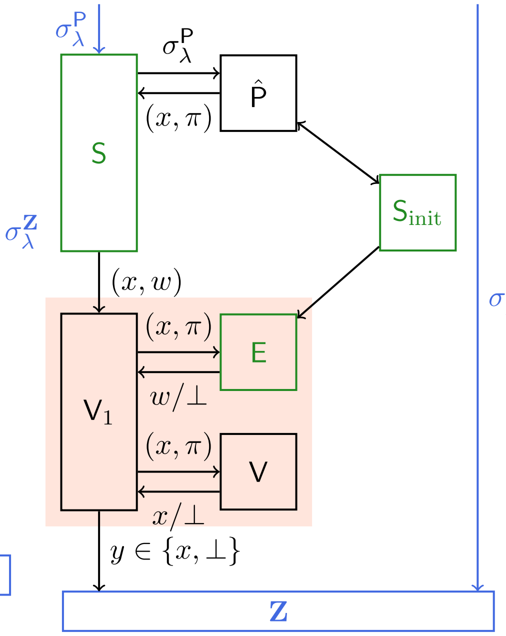

$\sigma_\lambda^{\mathsf{P}}$
$\sigma_\lambda^{\mathsf{P}}$
S
$\hat{\mathsf{P}}$
(x,π)
S_{init}
(x,w)
(x,π)
E
V_1
$w/\bot$
(x,π)
V
$x/\bot$
$y \in \{x,\bot\}$
$\sigma_\lambda^{\mathbf{Z}}$
Z

Figure 5.6: World_1

$$
\begin{aligned}
&= \Bigg| \Pr\big(\mathbf{Z}(x, \sigma_\lambda^{\mathbf{Z}}) = 1\big) \\
&\cdot \sum_{M \in \{\widetilde{\mathsf{V}}, \mathsf{V}_1\}} (-1)^{\delta(\mathsf{V})} \Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_M \langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \overset{\mathcal{F}}{\leftrightsquigarrow} M \rangle\Big) \\
&+ \Pr\big(\mathbf{Z}(\bot, \sigma_\lambda^{\mathbf{Z}}) = 1\big) \\
&\cdot \sum_{M \in \{\widetilde{\mathsf{V}}, \mathsf{V}_1\}} (-1)^{\delta(\mathsf{V})} \Pr\Big(y = \bot \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_M \langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \leftrightsquigarrow M \rangle\Big) \Bigg| \\
&\leq \left| \sum_{M \in \{\widetilde{\mathsf{V}}, \mathsf{V}_1\}} (-1)^{\delta(\mathsf{V})} \Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_M \langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \overset{\mathcal{F}}{\leftrightsquigarrow} M \rangle\Big) \right| \\
&+ \left| \sum_{M \in \{\widetilde{\mathsf{V}}, \mathsf{V}_1\}} (-1)^{\delta(\mathsf{V})} \Pr\Big(y = \bot \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_M \langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \overset{\mathcal{F}}{\leftrightsquigarrow} M \rangle\Big) \right| .
\end{aligned} \tag{5.55}
$$

We write down the probabilities of each protocol (dummy verifier $\widetilde{\mathsf{V}}$ from figure 5.5 and V_1 from figure 5.6) accepting to better visualize. In World_0, the probability of the protocol accepting is:

$$\Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\widetilde{\mathsf{V}}} \langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \overset{\mathcal{F}}{\leftrightsquigarrow} \widetilde{\mathsf{V}} \rangle\Big) \tag{5.56}$$

$$= \Pr\Big[y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), w' \leftarrow \mathsf{E}(x, \pi, H), y \leftarrow \mathcal{F}_{zk}^{\mathcal{R}}(x, w')\Big] \tag{5.57}$$

$$= \Pr\Big[w' \neq \bot \wedge (x, w') \in \mathcal{R} \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), w' \leftarrow \mathsf{E}(x, \pi, H)\Big] . \tag{5.58}$$

In World_1, the probability of the protocol accepting is:

$$\Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{V_1} \langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \leftrightsquigarrow V_1 \rangle\Big) \tag{5.59}$$

$$
\begin{aligned}
= \Pr\Big[y = x \wedge w' \neq \bot \Big| &H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), \\
&w' \leftarrow \mathsf{E}(x, \pi, H), y \leftarrow \mathsf{V}(x, \pi)\Big]
\end{aligned} \tag{5.60}
$$

$$\begin{aligned}
&= \Pr\Big[y = x \wedge w' \neq \bot \wedge (x, w') \in \mathcal{R} \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), \\
&\qquad\qquad w' \leftarrow \mathsf{E}(x, \pi, H), y \leftarrow \mathsf{V}(x, \pi)\Big] \\
&\quad + \Pr\Big[y = x \wedge w' \neq \bot \wedge (x, w') \notin \mathcal{R} \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), \\
&\qquad\qquad w' \leftarrow \mathsf{E}(x, \pi, H), y \leftarrow \mathsf{V}(x, \pi)\Big],
\end{aligned} \tag{5.61}$$

where in the last equality we just used the marginal probability expansion. Moreover, by online-extractability (see Definition 5.4.1), we know that for all possible advice $\sigma_\lambda^{\mathsf{P}}$:

$$\begin{aligned}
&\Pr\Big[y = x \wedge w' \neq \bot \wedge (x, w') \notin \mathcal{R} \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), \\
&\qquad\qquad w' \leftarrow \mathsf{E}(x, \pi, H), y \leftarrow \mathsf{V}(x, \pi)\Big] < \mathsf{negl}(\lambda),
\end{aligned} \tag{5.62}$$

which bounds eq. (5.58), the probability of the protocol accepting in World_0. Online-extractability also implies that

$$\begin{aligned}
&\Pr\Big[y = x \wedge w' \neq \bot \wedge (x, w') \in \mathcal{R} \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), \\
&\qquad\qquad w' \leftarrow \mathsf{E}(x, \pi, H), y \leftarrow \mathsf{V}(x, \pi)\Big] > 1 - \mathsf{negl}(\lambda),
\end{aligned} \tag{5.63}$$

thus bounding the first term on eq. (5.61), the probability of the protocol accepting in World_1. Note that we can bound the second term in eq. (5.61) by eq. (5.58), as the former adds one more restriction to the output. Therefore, both protocols $\widetilde{\mathsf{V}}$ and V_1 have similar probability of accepting

$$\begin{aligned}
&\Big|\Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\widetilde{\mathsf{V}}}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \overset{\mathcal{F}}{\leftrightsquigarrow} \widetilde{\mathsf{V}} \rangle\Big) \\
&\quad - \Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\mathsf{V}_1}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \leftrightsquigarrow \mathsf{V}_1 \rangle\Big)\Big|
\end{aligned} \tag{5.64}$$

$$\begin{aligned}
&\leq \Big|\Pr\Big[w' \neq \bot \wedge (x, w') \in \mathcal{R} \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), w' \leftarrow \mathsf{E}(x, \pi, H)\Big] \\
&\quad - \Pr\Big[y = x \wedge w' \neq \bot \wedge (x, w') \in \mathcal{R} \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), \\
&\qquad\qquad w' \leftarrow \mathsf{E}(x, \pi, H), y \leftarrow \mathsf{V}(x, \pi)\Big]\Big| + \mathsf{negl}(\lambda)
\end{aligned} \tag{5.65}$$

$$\leq \mathsf{negl}(\lambda). \tag{5.66}$$

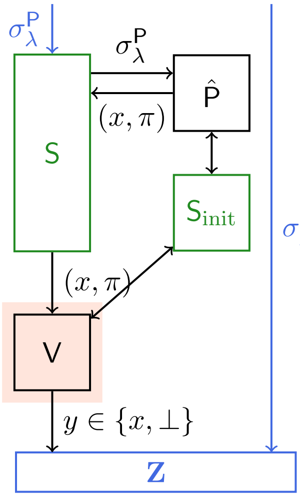
$\sigma_\lambda^{\mathsf{P}}$
$\sigma_\lambda^{\mathsf{P}}$
$\hat{\mathsf{P}}$
(x,π)
S
S_{init}
$\sigma_\lambda^{\mathbf{Z}}$
(x,π)
V
$y \in \{x, \bot\}$
$\mathbf{Z}$

Figure 5.7: World_2

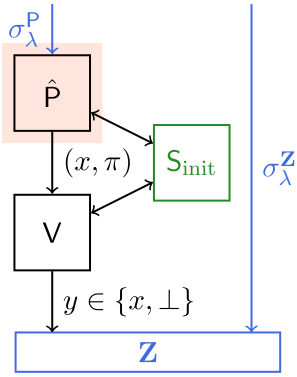
$\sigma_\lambda^{\mathsf{P}}$
$\hat{\mathsf{P}}$
S_{init}
(x,π)
$\sigma_\lambda^{\mathbf{Z}}$
V
$y \in \{x, \bot\}$
$\mathbf{Z}$

Figure 5.8: World_3

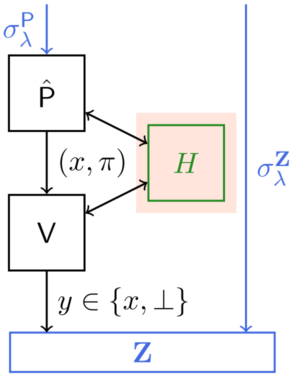

Figure 5.9: World_4

This bound in enough to show the similarity of World_0 and World_1 as

$$\begin{aligned}\Pr\Big(y = \bot \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_M\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \overset{\mathcal{F}}{\leftrightsquigarrow} M\rangle\Big)\\ = 1 - \Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_M\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \overset{\mathcal{F}}{\leftrightsquigarrow} M\rangle\Big), \quad (5.67)\end{aligned}$$

for both $M \in \{\widetilde{\mathsf{V}}, \mathsf{V}_1\}$.

We can also prove $\mathsf{World}_1 \approx \mathsf{World}_2$ using online extractability, as the verifier V_1 is only using the extractor to see if it does not abort. More precisely, following the same argument as before, and expanding the probability of accepting for the protocol from World_2:

$$\Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \leftrightsquigarrow \mathsf{V}\rangle\Big) \quad (5.68)$$

$$= \Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x,\pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), y \leftarrow \mathsf{V}(x,\pi)\Big) \quad (5.69)$$

$$\begin{aligned}&= \Pr\Big(y = x \wedge (x, w') \in \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x,\pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), y \leftarrow \mathsf{V}(x,\pi)\Big)\\ &\quad + \Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x,\pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), y \leftarrow \mathsf{V}(x,\pi)\Big).\end{aligned} \quad (5.70)$$

In order to prove the indistinguishability, we show that the probability of accepting is similar for both verifiers V_1 and V:

$$\begin{aligned}
&\Big|\Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\mathsf{V}_1}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \leftrightsquigarrow \mathsf{V}_1\rangle\Big) \\
&\quad - \Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, y \leftarrow \mathsf{OUT}_{\mathsf{V}}\langle \mathsf{Sim}_{\hat{\mathsf{P}}}(\sigma_\lambda^{\mathsf{P}}) \leftrightsquigarrow \mathsf{V}\rangle\Big)\Big|
\end{aligned} \tag{5.71}$$

$$\begin{aligned}
&= \Big| \Pr\Big[y = x \wedge w' \neq \bot \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), \\
&\qquad\qquad\qquad\qquad w' \leftarrow \mathsf{E}(x, \pi, H), \leftarrow \mathsf{V}(x, \pi)\Big]
\end{aligned} \tag{5.72}$$

$$\quad - \Pr\Big(y = x \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), y \leftarrow \mathsf{V}(x, \pi)\Big)\Big|$$

$$\begin{aligned}
&= \Pr\Big[y = x \wedge w' = \bot \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), \\
&\qquad\qquad\qquad\qquad w' \leftarrow \mathsf{E}(x, \pi, H), y \leftarrow \mathsf{V}(x, \pi)\Big]
\end{aligned} \tag{5.73}$$

$$\begin{aligned}
&= \Pr\Big[y = x \wedge w' = \bot \wedge (x, w') \notin \mathcal{R} \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \hat{\mathsf{P}}^H(\sigma_\lambda^{\mathsf{P}}), \\
&\qquad\qquad\qquad\qquad w' \leftarrow \mathsf{E}(x, \pi, H), y \leftarrow \mathsf{V}(x, \pi)\Big]
\end{aligned} \tag{5.74}$$

$$< \mathsf{negl}(\lambda)\,, \tag{5.75}$$

where in eq. (5.74) we used that $\{w' = \bot\} \subseteq \{(x, w') \notin \mathcal{R}\}$.

The rest of the relations are obvious since World_3 is just World_2 with a different routing of the messages (the simulator S is only redirecting the information). $\mathsf{World}_3 \approx \mathsf{World}_4$ is changing the S_{init} by the oracle that it is simulating H.

Case 3: malicious Berain. If the adversary corrupts the verifier $\mathcal{A} = \hat{\mathsf{V}}$, we will use the zero-knowledge property to construct the desired simulator.

Note that in Unruh's description of the adversary, the definition also encompasses the distinguisher, but by allowing adversaries that receive advice, it is equivalent to assuming an adversary that outputs a proof for the verifier. This is, the Unruh's simulator $\mathsf{Sim} = (\mathsf{S}_{\text{init}}, \mathsf{S}'_{\mathsf{P}})$ fulfills

$$\begin{aligned}
&\Big|\Pr\Big(\mathbf{Z}(y, \sigma_\lambda^{\mathbf{Z}}) = 1 \Big| H \leftarrow \mathsf{ROdist}, (x, \pi) \leftarrow \mathsf{P}^H(\sigma_\lambda^{\mathsf{P}}), y \leftarrow \hat{\mathsf{V}}(x, \pi, \sigma_\lambda^{\mathsf{V}})\Big) \\
&\quad - \Pr\Big(\mathbf{Z}(y, \sigma_\lambda^{\mathbf{Z}}) = 1 \Big| H \leftarrow \mathsf{S}_{\text{init}}, (x, \pi) \leftarrow \mathsf{S}'_{\mathsf{P}}(\sigma_\lambda^{\mathsf{P}}), y \leftarrow \hat{\mathsf{V}}(x, \pi, \sigma_\lambda^{\mathsf{V}})\Big)\Big| \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad < \mathsf{negl}(\lambda)\,.
\end{aligned} \tag{5.76}$$

Actually, the simulator S'_{P} is replacing both the ideal functionality $\mathcal{F}_{zk}^{\mathcal{R}}$ and the dummy verifier $\widetilde{\mathsf{P}}$ in the ideal world, i.e. aborts whenever $(x, w) \notin \mathcal{R}$ and runs $\mathsf{S}_{\mathsf{P}}(x, \sigma_\lambda^{\mathsf{V}})$ otherwise. However, we can easily modify this simulator to obtain the desired one in terms of the subsimulator S_{P}.

Protocol 21 Simulator $\mathsf{Sim}_{\hat{\mathsf{V}}} := (\mathsf{S}, \mathsf{S}_{\text{init}})$ for malicious verifier $\mathcal{A} = \hat{\mathsf{V}}$.

```
1: If $\mathsf{S}$ does not receive an abort $\perp$ message from the ideal functionality $\mathcal{F}_{zk}^{\mathcal{R}}$,
   redirects the input $x$ to the simulator $\mathsf{S}_{\mathsf{P}}$. Else, it aborts.
2: $\mathsf{S}$ receives $(x, \pi)$ from the simulator $\mathsf{S}_{\mathsf{P}}(x)$.
3: $\mathsf{S}$ sends $(x, \pi, \sigma_\lambda)$ to the verifier $\mathsf{V}$.
4: $\mathsf{S}$ redirects the output of the verifier $\mathsf{V}$ to the distinguisher $\mathbf{Z}$.
```

Let $\mathsf{Sim}_{\hat{\mathsf{V}}} := (\mathsf{S}, \mathsf{S}_{\text{init}})$ be the simulator from Protocol 21, since this is just a rewiring of Unruh's simulator, it is clear that $\mathbf{Z}$ cannot distinguish between the proofs provided by the verifier and the simulator, see the comparison between figures 5.10 and 5.11. □

5.5 Conclusion

In this chapter we presented a quantum oblivious transfer (OT) protocol in the stand-alone security model, which makes use of classical (non-interactive) zero-knowledge ((NI)ZK) to prove in which basis a quantum state is, without destroying it. Moreover, since our protocol inherits the security guarantees of the underlying (NI)ZK protocol, we can achieve a 2-message OT protocol, by considering Unruh's NIZK protocol (in the quantum random oracle model). Our contribution closes an open standing question, the round-optimality of *quantum* OT, which turns out to be equal to the round-optimality of classical OT.

The protocol has many moving parts, which open the door to questions. To name a few:

- Reducing entanglement: our protocols require the preparation of states representing a superposition of bit strings, and the application of a hash function h in superposition. For practical considerations, it would be great to see if we could get 2-message quantum OT protocols with single-qubit operations (or prove impossibility results).

- Universal composability: the model of security we are using allows sequential composability but not parallel composability. A priori, we expect our proof method to extend to a general composability frame-

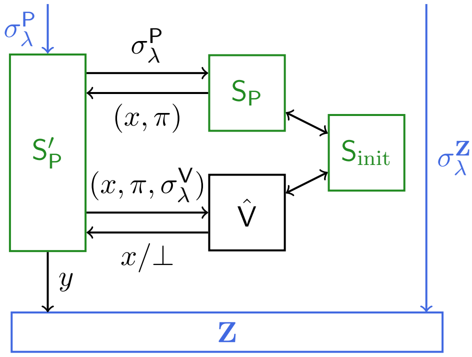

$\sigma_\lambda^{\mathsf{P}}$
$\sigma_\lambda^{\mathsf{P}}$
S_P
(x, π)
S'_P
$\mathsf{S}_{\mathrm{init}}$
$\sigma_\lambda^{\mathbf{Z}}$
$(x, \pi, \sigma_\lambda^{\mathsf{V}})$
$\hat{\mathsf{V}}$
$x/\bot$
y
$\mathbf{Z}$

Figure 5.10: Protocol for S'_P.

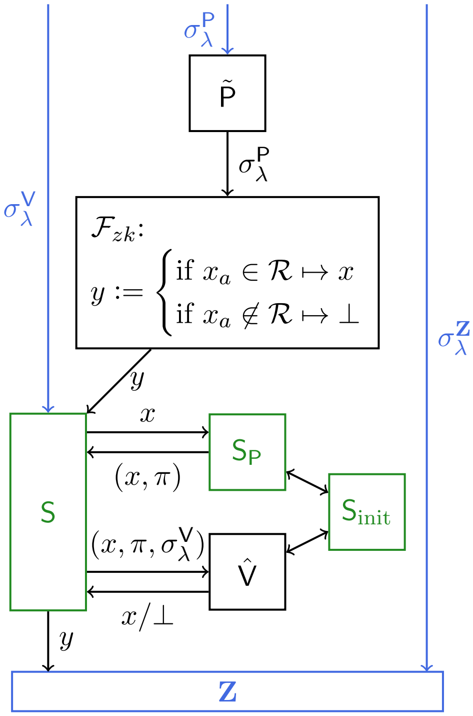

$\sigma_\lambda^{\mathsf{P}}$
$\tilde{\mathsf{P}}$
$\sigma_\lambda^{\mathsf{P}}$
$\sigma_\lambda^{\mathsf{V}}$
$\mathcal{F}_{zk}$:
$y := \begin{cases} \text{if } x_a \in \mathcal{R} \mapsto x \\ \text{if } x_a \notin \mathcal{R} \mapsto \bot \end{cases}$
$\sigma_\lambda^{\mathbf{Z}}$
y
x
S_P
(x, π)
S
$\mathsf{S}_{\mathrm{init}}$
$(x, \pi, \sigma_\lambda^{\mathsf{V}})$
$\hat{\mathsf{V}}$
$x/\bot$
y
$\mathbf{Z}$

Figure 5.11: Ideal world with simulator $\mathsf{Sim}_{\hat{\mathsf{V}}}$.

work as discussed in Section 2.2, but we also need to find ZK protocols secure in this stronger model of security.

- ZK for statistical security: While our approach states that we can get quantum OT with statistical security assuming the existence of a statistical ZK argument of knowledge (for unbounded verifier/sender) or a ZK proof of knowledge (for unbounded prover/receiver), it is important to check that such protocols exist (for now the protocols we analyse only bring computational security, which results in a computationally secure OT).

- Even weaker assumptions: the hash function needs to be hiding (our actual assumption is actually slightly weaker), but we do not know if we can reduce this assumption to use only one-way functions.

- Reducing communication in the plain model: while our approach can get us to the optimal round-complexity (2 messages), such optimal complexity cannot be obtained in the plain-model, at least in a composable framework. It would be interesting to study the minimum number of rounds in the plain-model.

Chapter 6

Background on Quantum Pseudorandomness

This chapter provides an overview of quantum pseudorandomness, useful for understanding Chapters 7 and 8; it is original to the thesis and has not appeared in prior publications.

We begin in Section 6.2 by defining core pseudorandom primitives and examining different aspects of their definitions. Finally, in Section 6.3, we provide an overview of the established relations among these primitives.

6.1 Introduction

As mentioned in the introduction, cryptography fundamentally relies on the existence of problems that are difficult to solve but easy to verify. In classical cryptography, such problems are abstracted through one-way functions, which serve as a central building block for most protocols. When we transition to the post-quantum setting, i.e. where adversaries may possess quantum computers, these one-way functions must remain hard to invert even for quantum solvers.

However, once we consider honest parties who also have access to quantum resources, the long-held computational hardness requirements must be revisited. For example, the BB84 protocol demonstrates that key exchange can be achieved unconditionally when parties are allowed to communicate quantum states (albeit assuming the existence of a classically secure authenticated channel). More recently, Kretschmer [Kre21] showed that quantum pseudorandomness, and hence one-wayness, is possible even if $\mathsf{BQP} = \mathsf{QMA}$, the quantum analogue of $\mathsf{P} = \mathsf{NP}$. This stands in stark contrast to classical cryptography, where the existence of one-way functions requires that $\mathsf{P} \neq \mathsf{NP}$.

This divergence between classical and quantum cryptography has fuelled significant research interest, leading to a proliferation of new and often complex cryptographic primitives. While this corner of cryptography advances rapidly, and some of today's results may soon be refined or superseded, I hope the overview included here will provide a useful foundation for understanding the later chapters of the thesis, and serve as a guide for those entering the field.

6.2 Defining primitives

Classically, a distribution is considered pseudorandom if samples drawn from it are computationally indistinguishable from samples drawn from a uniform distribution. The quantum analogue, introduced by Ji, Liu and Song [JLS18], shifts the framework from classical bit-strings to quantum states, with randomness measured against the Haar distribution, the unique uniform distribution over the space of states.

We can differentiate between *decision* and *search* primitives. Of the first variant are for example pseudorandom state generators (PRSG), defined as classical-input, quantum-output algorithms that produce states indistinguishable from Haar-random states given a random seed, even when the distinguisher is provided polynomially-many copies of the states. This concept naturally extends to other decision-based quantum primitives. For instance, pseudorandom unitaries (PRUs), which are defined replacing state outputs by unitaries, act as a quantum counterpart to classical pseudorandom functions. Meanwhile, efficiently generatable, statistically far but computationally indistinguishable pairs of states (EFIs) serve as a quantum analogue to classical EFIDs. Search primitives, on the other hand, include one-way state generators (OWSG), where the challenge lies in finding a preimage of a quantum state generated by an algorithm with classical seeds.

Unlike their classical counterparts, quantum pseudorandom primitives do not inherently rely on the assumption that verifying is easier than solving, which suggests that these quantum constructions may operate under weaker computational assumptions than classical one-way functions. Despite this, they enable non-trivial cryptographic protocols, including quantum bit commitment and quantum multi-party computation [AQY22, MY22, HMY23b].

Another critical distinction between classical and quantum cryptography lies in the relationship between primitives. In the classical world, pseudorandom generators, pseudorandom functions and one-way functions are equivalent. In the quantum realm, however, these equivalences break down and the precise relationships between these primitives remain unclear. To understand the known nuances in the hierarchy of primitives, we will first develop an intuition about the key features distinguishing them and examine how their differing properties enable specific cryptographic applications.

Formal definitions. We include first the formal definitions of the aforementioned quantum primitives.

Definition 6.2.1 (PRSG)**.** *Let $\lambda \in \mathbb{N}$, ℓ be a polynomial and $\mathcal{K} \subseteq \{0,1\}^\lambda$. A* QPT *algorithm G is a* pseudorandom quantum state generator *(PRSG) if:*

i. ***State Generation:*** *There exist $\ell(\lambda)$-qubit pure states $\{|\varphi_k\rangle\}_{k\in\mathcal{K}}$ such that*

$$G(1^\lambda, k) = |\varphi_k\rangle\langle\varphi_k| \quad \textit{for all } k \in \mathcal{K}\,. \tag{6.1}$$

ii. ***Pseudorandomness:*** *For any* QPT *observer $\mathcal{A}$ and all polynomials* *t, we have:*

$$\left|\Pr_{k\leftarrow\{0,1\}^\lambda}\left(\mathcal{A}\left(1^\lambda,|\varphi_k\rangle^{\otimes t(\lambda)}\right)=1\right)\right.$$
$$\left.-\Pr_{|\nu\rangle\leftarrow\sigma_{\ell(\lambda)}}\left(\mathcal{A}\left(1^\lambda,|\nu\rangle^{\otimes t(\lambda)}\right)=1\right)\right|\leq \mathsf{negl}(\lambda)\,, \quad (6.2)$$

where σ_n is the Haar distribution over n-qubit states $\mathbb{S}(2^n)$.

Sometimes authors refer to PRS instead, dropping the generator, by focusing on the discrete family of states $\{|\varphi_k\rangle\}_{k\in\mathcal{K}}$, but these still need to be efficiently generatable. PRSGs generalize classical pseudorandom generators (PRG) by giving a family of quantum states (instead of a family of bit-strings) that is computationally indistinguishable from truly Haar-random states (instead of indistinguishable from uniformly sampling bit-strings). Note that a uniform sampling of bit-strings is already sampling from a discrete distribution, which is not the case for the space of quantum states, and it is often required in the definition of PRGs that the output space is larger than the input, i.e. $\ell(\lambda)>\lambda$, which is not included in the PRSG definition.

Definition 6.2.2 (EFI)**.** *We call $\{\rho_0(\lambda),\rho_1(\lambda)\}_{\lambda\in\mathbb{N}}$ an* efficiently samplable, statistically far but computationally indistinguishable *(EFI) family of pairs of states if:*

i. ***Efficient generation:*** *There exists a* QPT *algorithm* EFI *such that*

$$\mathsf{EFI}(1^\lambda,b)=\rho_b(\lambda)\,. \quad (6.3)$$

ii. ***Statistically far:*** *There exists a polynomial p such that*

$$\|\rho_0(\lambda)-\rho_1(\lambda)\|_1\geq 1/p(\lambda)\,. \quad (6.4)$$

iii. ***Computational indistinguishability:*** *For any* QPT *observer $\mathcal{A}$ we have:*

$$|\Pr(\mathcal{A}(\rho_0(\lambda))=1)-\Pr(\mathcal{A}(\rho_1(\lambda))=1)|\leq\mathsf{negl}(\lambda)\,. \quad (6.5)$$

EFIs can be understood as a quantum generalization of EFIDs, a family of pairs of distributions that are statistically far yet computationally indistinguishable, where the distributions are now allowed to have quantum correlations. Note that the observer only has access to one copy of each state, as opposed to a polynomial number of copies in the security definition of PRSGs 6.2. This is because in the security requirement of EFIs we are trying to distinguish two *fixed* states and by parallel repetition if single-copy distinguishability holds then multi-copy distinguishability holds immediately.

Definition 6.2.3 (OWSG)**.** *Let* $\lambda \in \mathbb{N}$ *and* ℓ *be a polynomial. We say that a tuple of algorithms* $(\mathsf{KeyGen}, \mathsf{StateGen}, \mathsf{Ver})$ *is a* one-way state generator *(OWSG) if the algorithms satisfy the following specifications:*

i. ***Key Generation:*** $\mathsf{KeyGen}(1^\lambda) \to k$*: is a* QPT *algorithm.*

ii. ***State Generation:*** $\mathsf{StateGen}(1^\lambda, k) = \rho_k$*, for some* $\ell(\lambda)$*-qubit mixed state* $\rho_k \in \mathsf{MS}_{\ell(\lambda)}$*, is a* QPT *algorithm.*

iii. ***Verification:*** $\mathsf{Ver}(k, \rho) \in \{\perp, \top\}$ *is a (possibly unbounded) quantum algorithm.*

Further, they satisfy the following properties:

1. ***Correctness:*** *Outputs of the samplers* $(\mathsf{KeyGen}, \mathsf{StateGen})$ *pass the verification with overwhelming probability, i.e.,*

$$\Pr_{\substack{k \leftarrow \mathsf{KeyGen} \\ \rho_k \leftarrow \mathsf{StateGen}(k)}} (\mathsf{Ver}(k, \rho_k) = \top) \geq 1 - \mathsf{negl}(\lambda)\,. \tag{6.6}$$

2. ***Security:*** *For every* QPT *observer* $\mathcal{A}$*, and any polynomial* t*:*

$$\Pr_{\substack{k \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \rho_k \leftarrow \mathsf{StateGen}(1^\lambda, k) \\ k' \leftarrow \mathcal{A}(\rho_k^{\otimes t(\lambda)})}} \big(\mathsf{Ver}(k', \rho_k) = \top\big) \leq \mathsf{negl}(\lambda)\,. \tag{6.7}$$

OWSGs are quantum generalization of one-way functions (OWF), as the $\mathsf{StateGen}$ algorithm efficiently computes a state associated with every key $k \in \mathrm{Im}(\mathsf{KeyGen})$, but by the security definition in eq. (6.7) no efficient algorithm can find a pre-image that passes verification with non-negligible probability.

To understand why the above definitions are chosen as quantum generalizations, we have to look at different properties that are implicit. We will now go through the nuances.

Number of copies. Note that in the definitions of PRSGs and OWSGs we define security in terms of distinguishers that have access to a polynomial number of copies of the states. In contrast to the classical case, where the ciphertexts can be copied and there is no difference between how many copies the observer is provided, quantum states cannot be *cloned*, thus it is essential to state how many copies the observer has access to, where the more copies the higher the security guarantee. However, for some cryptographic tasks it is not necessary to ask for such a stringent security definition, we then instead consider t-copy secure PRSGs, also called t-PRSGs.

Computationally bounded distinguisher. An essential part of Definitions 6.2.1 and 6.2.3 is requiring security with respect to computationally *bounded* distinguishers, QPT algorithms $\mathcal{A}$. If $\mathcal{A}$ is instead a computationally *unbounded* algorithm we would have an information-theoretical primitive. A closely linked concept is that of state t-designs, a discrete set of efficiently computable states that are information-theoretically indistinguishable from t-copies of a Haar-random state, but with no expansion requirement, i.e. $\ell(\lambda) = \lambda$. Although t-designs are sometimes considered a pseudorandom primitive, it does not refer to the same pseudorandomness notion as in cryptography, where it must explicitly address the computational boundedness of the distinguisher. Moreover, state t-designs without any expansion requirement, i.e. the key size being similar to the number of states $\ell(\lambda) = \lambda$, do exist unconditionally [DCEL09, HMMH$^+$23, OSP23] (and so do non-expanding t-PRS).

Output length. As we just hinted at, the existence of the primitives is conditional on the number of qubits $\ell(\lambda)$ in the output, in contrast to the classical case. In PRGs pseudorandom bit-strings are pseudorandom bitwise, so a smaller pseudorandom bit-string can be obtained by ignoring part of the string. There are also ways of amplifying the output length of a PRG [Gol06]. However, length shrinking or expansion does not seem to trivially hold quantumly. For example, PRSGs of output length $\ell(\lambda) = c \cdot \log(\lambda)$ for $c < 1$ exist unconditionally, while this is not the case for $c > 1$ or for linear length PRSGs [BS20b]. Note that although tempting, trivially doing tomography to learn the states as a mean of distinguishing them from Haar-random ones does not work, as the error induced by tomography is too high for non-negligibly distinguishing between them. Although there is currently no known way of length expansion of PRSGs preserving the security for a polynomial number of copies of the states, Gunn, Ju, Ma and Zhandry showed that $\omega(\lambda)$-length PRSGs imply 1-PRSGs with larger length [GJMZ23]. In the rest of this thesis we will distinguish between the two extreme cases of computational PRSGs, we will refer to PRSGs whose output length is logarithmic (with lower constant $c > 1$) in the security parameter $\ell(\lambda) = \Theta(\log(\lambda))$ as short-PRSG, and keep the term PRSGs for states whose length is super-logarithmic $\ell(\lambda) = \omega(\log(\lambda))$. Note that Kretschmer's separation [Kre21] is only between PRSGs of super-logarithmic length and OWF, and does not address short-PRSGs.

Pseudo-determinism. Although the tomography error is still too high to construct algorithms with deterministic output, logarithmic output length PRSGs do admit efficient tomography. Motivated by this Ananth, Lin and Yuen [ALY24] proposed a way of building *quantum* PRGs from measuring, and rounding, the output of a short-PRSG. By the inherent probabilistic nature of the measurements a deterministic outcome is not possible, which lead to the introduction of *pseudo-deterministic* cryptographic primitives. Let us define pseudo-deterministic PRGs to further illustrate the properties of these primitives.

Definition 6.2.4 (PD-PRGs)**.** *Let $\lambda \in \mathbb{N}$ and let m, ℓ be two polynomials. A* QPT *algorithm G is a* quantum pseudo-deterministic pseudorandom generator *(PD-PRG) if the following conditions hold:*

i. ***Pseudo-determinism:*** *For every $\lambda \in \mathbb{N}$, there exists a set $\mathcal{K}_\lambda \subset \{0,1\}^{m(\lambda)}$ with*

(a) $\Pr_{x \leftarrow \{0,1\}^{m(\lambda)}}(x \in \mathcal{K}_\lambda) \geq 1 - \mathsf{negl}(\lambda)$.

(b) For any $x \in \mathcal{K}_\lambda$, it holds that

$$\max_{y \in \{0,1\}^{\ell(\lambda)}} \Pr\Big(y = G(1^\lambda, x)\Big) \geq 1 - \mathsf{negl}(\lambda)\,, \tag{6.8}$$

where the probability is over the randomness of G.

ii. ***Security:*** *For every* QPT *inverter $\mathcal{A}$:*

$$\Bigg| \Pr_{k \leftarrow \{0,1\}^{m(\lambda)}} \Big(\mathcal{A}(1^\lambda, G(1^\lambda, k)) = 1\Big) - \Pr_{r \leftarrow \{0,1\}^{\ell(\lambda)}} \Big(\mathcal{A}(1^\lambda, r) = 1\Big) \Bigg| \leq \mathsf{negl}(\lambda)\,. \tag{6.9}$$

where the probability is over the randomness of G and $\mathcal{A}$.

In a way, these primitives are not necessarily functions as they could act arbitrarily on inputs outside the key space $\mathcal{K}_\lambda$. In [ALY24] the authors showed that short-PRSGs imply a variant of PD-PRGs where determinism only holds with inverse-polynomial probability, denoted PD-PRGs$^{\mathsf{poly}}$, which they proved to be enough to construct EFIs and pseudo-encryption. However, they were not able to build PD-PRGs from short-PRSGs, which would imply black-box length expansion of PRSGs. It is an interesting open question to see under what circumstances pseudo-deterministic quantum primitives are enough for constructing cryptographic protocols.

Hybrid cryptography. Pseudo-randomness is a useful tool for analysing the possibilities of cryptography with quantum computational capabilities but limited to classical communication, sometimes referred to as QCCC. One-way puzzles, a variant of one-way functions where the sampling of the image and the key are done jointly, imply almost all QCCC primitives as shown by Khurana and Tomer [KT24]. We will not include their definition, as we will not directly work with them throughout the rest of the thesis, but just mention that they are implied by PRSGs.

Verification. In contrast to classical one-way functions, one-way state generators are defined with respect to a verification algorithm. Recall that classically given $f(k)$ we intend to find a key k, which can be verified by simply given a candidate k', by computing $f(k')$ and checking if it equals $f(k)$. For quantum states, however, this check is non-trivial. If the OWSG outputs were pure $\rho_k = |\phi_k\rangle\langle\phi_k|$, the verification algorithm could, for instance, measure the output with $\{|\phi_k\rangle\langle\phi_k|, I - |\phi_k\rangle\langle\phi_k|\}$. If the OWSG outputs were mixed, we could similarly check the trace distance with respect to the ideal expected output, but this procedure is not efficient. Different works in the literature consider different definitions for OWSG; with either pure or mixed outputs, and with either efficient or inefficient verification procedures.

The distinction of efficient and inefficient verification is not restricted to quantum output primitives. Recently, Chung, Golding and Gray showed that there is a quantum algorithm with respect to which inefficiently verifiable one-way puzzles exist whilst efficiently verifiable ones do not [CGG24]. The proof relies on the fact that efficiently-verifiable one-way puzzles can be broken with a QCMA oracle, but as mentioned earlier Kretschmer showed that PRSGs (and thus inefficiently verifiable one-way puzzles) can exist even if $\mathsf{BQP} = \mathsf{QCMA}$ [Kre21].

Pseudoresources. Haar-random states are maximally entangled with high probability, and similarly states in any PRS are highly entangled on average [JLS24]. Instead of randomness, Gulão and Elkouss [GE24] showed that *pseudoentangled* states, quantum states computationally indistinguishable from maximally entangled states, imply EFIs. Recently, Grilo and Yángüez [GY25] showed that this idea can be generalized to many types of quantum *pseudoresources*.

Primitive construction. The first candidate construction for PRSGs was given by Ji, Li and Song [JLS24] from classical pseudorandom functions (PRFs); the construction consists of taking a random superposition of computational basis states, with the phase being determined by the PRF, where the inputs to the PRF are used for the superposition and the pseudorandom states are keyed with the same keys as the PRFs. Formally, let $\lambda \in \mathbb{N}$ and let $F\colon \mathcal{K} \times \mathcal{X} \to \mathcal{X}$ be a PRF with key space $\mathcal{K} \subseteq \{0,1\}^\lambda$ and input space $\mathcal{X} \subseteq \{0,1\}^\lambda$, then the family of states $\{|\varphi_k\rangle\}_{k\in\mathcal{K}}$ defined as

$$|\varphi_k\rangle := \frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x\in\mathcal{X}} \omega^{F_k(x)} |x\rangle \,, \tag{6.10}$$

forms a PRS, where ω is a $|\mathcal{X}| - 1$ root of unity. A later construction by Brakerski and Shmueli [BS20b] simplified the above superposition by only considering bit-flips as phase. Most other pseudorandom primitives are implied by PRSG, and this is definitely the case for OWSG [MY24a, CGG$^+$25] and EFIs [MY24a].

Recently there have been some advancements in the meta-complexity of quantum pseudorandomness. To name a few examples, Morimae, Nehoran and Yamakawa [MNY24] proved that quantum auxiliary-input EFIs exist unconditionally, and concurrent works by Cavalar, Golding, Gray and Hall [CGGH25] and the one of Hiroka and Morimae [HM24] based the existence of one-way puzzles, and hence EFIs, on the quantum average-hardness of the GapK problem. Finally, PRSGs were shown to exist conditional on the hardness of a new inherently quantum hardness assumption, the so-called Hamiltonian Phase State problem, by Bostanci, Haferkamp, Hangleiter and Poremba [BHHP24].

Protocols based on pseudorandom primitives. Quantum pseudorandomness has found many applications since its introduction. The definition of PRSGs naturally gives rise to private-key quantum money [JLS24], later

Morimae and Yamakawa [MY22] showed that non-interactive statistically-binding quantum commitments (which were proven by Brakerski, Canetti and Qian [BCQ23] and Yan [Yan22] to be equivalent to EFIs) can be constructed from 1-PRSGs. However, constructing quantum cryptographic tasks from PRSGs alone proved quite difficult, given the length rigidity and entanglement constrains we mentioned in the previous paragraph, which led to the introduction of *pseudorandom function-like state generators* (PRFSGs) by Ananth, Qian and Yuen [AQY22]. We include here the definition for clarity.

Definition 6.2.5 (PRFSG)**.** *Let $\lambda \in \mathbb{N}$, ℓ and d be polynomials and $\mathcal{X}, \mathcal{K} \subseteq \{0,1\}^\lambda$. A* QPT *algorithm G is a (selectively-secure)* pseudorandom quantum function-like state generator *(PRFSG) if:*

i. ***State Generation:*** *For every polynomial s, there exist pure states $\{|\varphi_{x_j,k}\rangle\}_{j\in[s(\lambda)],x\in\mathcal{X}^{d(\lambda)},k\in\mathcal{K}}$ of size $\ell(\lambda)$, such that:*

$$G(1^\lambda, k, x_j) = |\varphi_{x_j,k}\rangle\langle\varphi_{x_j,k}| \quad \textit{for all } j \in [s(\lambda)], x_j \in \mathcal{X}^{d(\lambda)}, k \in \mathcal{K}\,. \tag{6.11}$$

ii. ***Pseudorandomness:*** *For any* QPT *observer $\mathcal{A}$ and all polynomials t and family of indices $\{x_1, \ldots, x_{s(\lambda)}\} \subseteq \{0,1\}^{d(\lambda)}$, we have:*

$$\begin{aligned}\Bigg| &\Pr_{k\leftarrow\{0,1\}^\lambda}\left(\mathcal{A}_\lambda(x_1,\ldots,x_s,|\varphi_{x_1,k}\rangle^{\otimes t},\ldots,|\varphi_{x_s,k}\rangle^{\otimes t}) = 1\right) \\ &- \Pr_{|\nu_1\rangle,\ldots,|\nu_s\rangle\leftarrow\sigma_{\ell(\lambda)}}\left(\mathcal{A}_\lambda(x_1,\ldots,x_s,|\nu_1\rangle^{\otimes t},\ldots,|\nu_s\rangle^{\otimes t}) = 1\right)\Bigg| \leq \mathsf{negl}(\lambda)\,,\end{aligned} \tag{6.12}$$

where s and t implicitly depend on the security parameter λ.

While PRSGs output a single state per key $k \in \mathcal{K}$, we can think of PRFSGs as outputting a tensor product of pseudorandom states. Note also that given a PRFSG G, the family of states $G(1^\lambda, k, 0^{d(\lambda)})$ forms trivially a PRS family. Although this new perspective allowed the authors to construct EFIs and quantum pseudo one-time pads from $d(\lambda) = \Theta(\log(\lambda))$ and $\ell(\lambda) = \omega(\log(\lambda))$ PRFSG, it turns out that for these particular input-output parameters PRFSGs can be constructed from PRSGs [AQY22]. As a final point, Kretschmer's oracle separation [Kre21] can also be extended, such that PRFSGs are a potentially weaker assumption than OWF.

Query models. The final property to discuss is the access models to these primitives. Firstly, the previously stated PRFSG was defined in terms of *selective security*, that is, the classical inputs of the distinguisher to the PRFSG are determined ahead of time. The primitives obtained from such definition will necessarily also be selectively secure. Ananth, Gulati, Qian and Yuen [AGQY22] later constructed also *adaptively secure* PRFSGs from PRFs, that is, the distinguisher has access to an oracle that on input $x \in \{0,1\}^{d(\lambda)}$ either outputs a PRSFG output state $|\varphi_{x,k}\rangle$ or a Haar-random state. For the particular instantiation of $d(\lambda) = \Theta(\log(\lambda))$, selectively and adaptively secure PRFSGs are equivalent. Finally, we need also consider to define the quantum pseudorandom primitives with either having classical or quantum access to these oracle.

6.3 Relations between primitives

The study of pseudorandom quantum primitives reveals a landscape of closely related yet subtly distinct definitions, often with slight variations in security requirements or functionalities. While these primitives share conceptual similarities with their classical counterparts, their relative strengths are not immediately clear. Although an updated map of these connections is maintained by Or Sattath [Sat], this section provides a concise overview, focusing on key implications and hierarchy results among the most prominent primitives.

The graph in figure 6.1 is drawn to suggest that the primitives on top are stronger requirements than the primitives on the bottom.

To remark a few interesting points in relation to the properties discussed in Section 6.2, we see that the input/output dimension of the primitives is crucial. On the one hand, logarithmic input PRFSGs are known to be equivalent to PRSGs, and hence separated from OWF via Kretschmer's oracle, but such results are not known for super-logarithmic input PRFSGs. On the other hand, PRSGs with different output lengths are also separated, suggesting that short-PRSGs are a stronger primitive than PRSGs.

Given that short-PRSGs are not known to be separated from OWF, while they are sufficient to give pseudo-deterministic variants of classical pseudorandomness, there might not be any clear advantage in considering them as primitives (if we wanted a strong primitive, we could just consider the well-studied one-way functions), thus we believe it is important to discuss other possibly weaker but still interesting pseudorandom primitives. EFIs

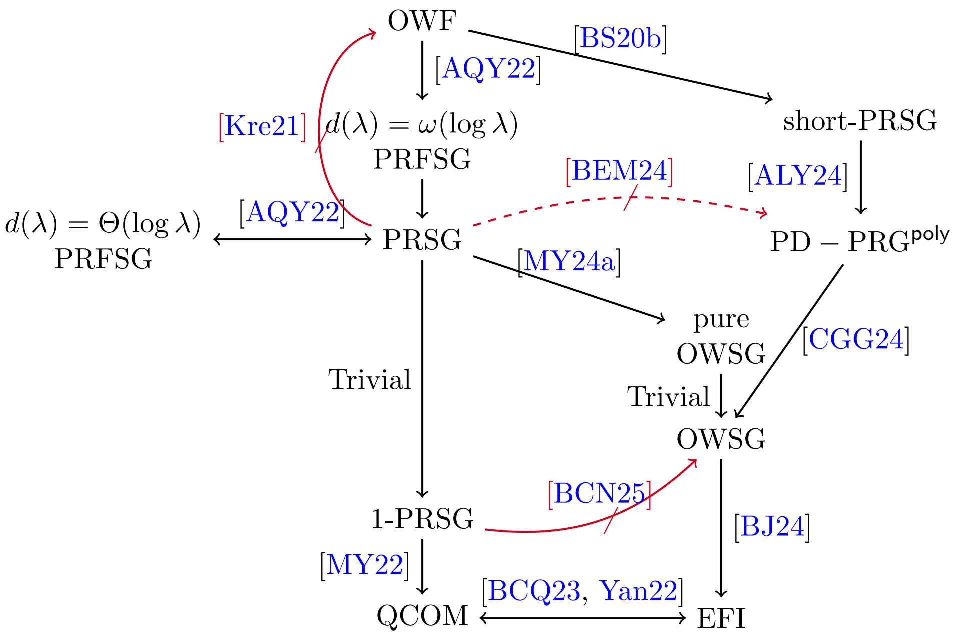

Figure 6.1: Map of relations between primitives. Arrows $A \to B$ denote that primitive B can be constructed from primitive A, arrows with a cross $A \nrightarrow B$ denote that there is an oracle world were primitive A exists but primitive B does not. The dashed arrow represents the relation discovered in an article in this thesis.

are the primary candidate for centrality in quantum pseudorandomness; not only are they equivalent to many other tasks such as quantum commitments or quantum multi-party computation, but they are also implied both by all the other quantum pseudorandom primitives and QCCC ones. Interestingly, 1-copy PRSGs, at the beginning believed to be quite weak primitives, are enough to construct EFIs.

The one-way state generators (OWSG) included in figure 6.1 are assumed to have efficient verification algorithms, as these are probably more useful. Given the separation between 1-PRSGs and OWSG, they seem to be of in-between hardness; strong enough to construct commitments but weaker than classical one-wayness. For example, whilst short-PRSGs and PRFSGs exit, defining short as the output length being of $\ell(\lambda) = \Theta(\log \lambda)$ qubits (for some lower constant $c > 1$), Hhan, Morimae and Yamakawa recently proved that short-OWSG with efficient verification algorithm and short-EFIs do not exist [HMY23a]. However, as discussed earlier, admitting either pure or mixed outputs and efficient or inefficient verification algorithms provides a set of incomparable OWSGs, thus it is important to pay attention to which type of primitive each article discusses. Interestingly, if we admit inefficient verification algorithms, Batra and Jain [BJ24] proved that OWSGs with mixed state outputs are equivalent to EFIs. Moreover, the number of copies given to the distinguisher also plays a role, as they proved that $O(\lambda/\log\lambda)$-copy OWSG are equivalent to OWSG. This is actually the optimal number of copies necessary for the existence of inefficiently verifiable OWSGs as was proved by Cavalar, Goldin, Gray, Hall, Liu and Pelecanos [CGG$^+$25].

Chapter 7

No shrinking of Quantum Pseudorandomness

This chapter is based on the article *Quantum Pseudorandomness Cannot Be Shrunk In a Black-Box Way* [BEM24], and is joint work with Samuel Bouaziz-Ermann.

A conclusion section has been added and remaining sections have been edited for style and typographical consistency.

In this chapter we prove that, relative to a quantum oracle, the existence of quantum pseudorandom state generators (PRSGs) does not guarantee the existence of pseudo-deterministic one-way functions (PD-OWFs), whose existence does imply PRSGs. Therefore, the output length of a PRSG cannot the shrunk in a black-box way.

7.1 Introduction

In classical cryptography, computational pseudorandomness generated by pseudorandom generators (PRGs) serves as a central resource. This primitive can be used for many applications such as commitments [Nao91], digital signatures [Rom90], and symmetric-key encryption. Furthermore, the existence of PRGs is necessary for the existence of almost all cryptographic primitives with computational security, including one-way functions (OWFs) [HILL99].

As we mentioned in Chapter 6, when we treat the world as operating under the laws of quantum mechanics, however, the notion of pseudorandomness must be revisited. Ji, Liu, and Song [JLS18] proposed the first inherently quantum pseudorandom primitive, an analogue to PRGs, called *Pseudorandom Quantum States* (PRSs) that consists of a family of polynomial size keyed-states $\{|\phi_k\rangle\}_{k\in\mathcal{K}}$ such that no quantum polynomial-time algorithm can distinguish between a polynomial number of copies of a randomly sampled element from the PRSs family or a polynomial number of copies of a Haar-random state. We sometimes instead talk about the generation algorithm that given a classical input key $k \in \mathcal{K}$ outputs a state $|\phi_k\rangle$ from the PRSs family, called *Pseudorandom Quantum State Generator* (PRSG). Quantum pseudorandomness has been shown to be useful to construct many (quantum) cryptographic primitives: Public key encryption with quantum keys [BGH$^+$23], quantum digital signatures [MY22], pseudo one-time pad encryption schemes [AQY22], statistically binding and computationally hiding commitments [AQY22, MY24a, KT24] and quantum computational zero knowledge proofs [BCQ23]. Such broad interest derives probably from the fact that PRSGs can be constructed from OWFs [JLS18] (and thus PRGs), but there are oracle separations found between OWFs and PRSGs [Kre21, KQST23], which makes them a potentially weaker building block for quantum cryptography, with a purely quantum description.

However, the landscape of quantum cryptographic pseudorandomness is quite different from the classical one. To begin with, the lengths of the primitives do not seem to be trivially interchangeable. Security proofs based on classical pseudorandomness can rely on *cleverly composing* the PRG with itself [Gol06, Section 3.3.2], which allows stretching the output length of the final pseudorandom string. However, no analogous intuitive construction is possible with PRSGs, because here the input is a classical string and the output is a quantum state. Moreover, although classically we can trivially use parts of the pseudorandom string independently, this is not possible quantumly. Discarding half of a quantum state could lead to a maximally mixed state, making it easy to distinguish the output the PRSGs output state from a Haar-random state, which means that pseudorandomness is not a property respected by taking subsets of the registers. Given that it is unclear how quantum pseudorandomness behaves with different output sizes, we start by considering the shortest possible cryptographic PRS, we refer to *short* quantum pseudorandomness when the output length is logarithmic in the security parameter. Similarly, we refer to *long* quantum pseudorandomness when the output length is super logarithmic, although we occasionally drop long as they are usually considered the standard definition. Several authors [BS20b, AQY22, BB21] have asked the following question

What is the relation between short and long PRSGs?

Our contributions. In this chapter we partially answer the above question; we show that there exists a quantum oracle relative to which PRSGs exist but short-PRSGs do not.

At the heart of our separation lies a new classical input/output but quantum algorithm that behaves *pseudo-deterministically.* Although a quantum state can encode arbitrary classical information, this information is not necessarily *accessible* for an observer, since measurements can incur errors. Pseudo-deterministic variants of classical cryptographic primitives were first introduced by Ananth, Lin and Yuen [ALY24], who defined a quantum polynomial-time algorithm that outputs a pseudorandom string on a *fraction* of the input keys, called pseudo-deterministic pseudorandom generators (PD-PRGs), and showed that these are enough to build short-PRSGs. Later Barhoush, Behera, Ozer, Salvail and Sattath [BBO$^+$24] introduced the analogous pseudo-deterministic one-way functions (PD-OWF). They showed that PD-OWFs can be built from PD-PRGs, and thus also from short-PRSGs.

In this chapter we show that Kretschmer's original oracle [Kre21] not only implies that one-way functions do not exist, but also that none of the pseudo-deterministic variants exist either. Kretschmer's oracle consist of two oracles: the first one is a quantum oracle, a keyed family of families of random unitaries, and hence can be interpreted as a quantum version of a random function, and the second part is a classical oracle that ensures $\mathsf{promiseBQP} = \mathsf{promiseQMA}$. It is the former part of the oracle that gives the existence of PRSGs, and here we show that if $\mathsf{promiseBQP} = \mathsf{promiseQMA}$, then we can break any construction of PD-OWFs. In particular, we prove the following result.

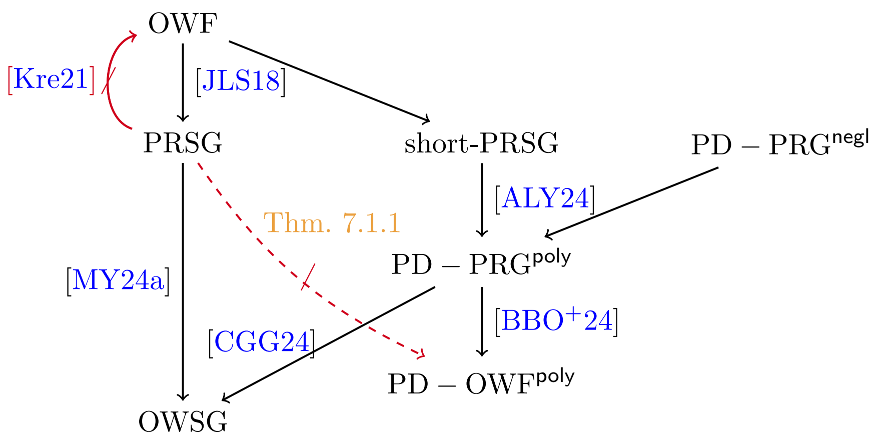
OWF
[Kre21]
[JLS18]
PRSG
short-PRSG
PD – PRG$^{\mathsf{negl}}$
[ALY24]
Thm. 7.1.1
PD – PRG$^{\mathsf{poly}}$
[MY24a]
[BBO$^+$24]
[CGG24]
PD – OWF$^{\mathsf{poly}}$
OWSG

Figure 7.1: Zoomed map of relations between primitives, for a general one see figure 6.1. Arrows $A \to B$ denote that primitive B can be constructed from primitive A, arrows with a cross $A \not\to B$ denote that there is an oracle world were primitive A exists but primitive B does not. The dashed arrow represents the relation discovered in the article that comprises this chapter.

Theorem 7.1.1 (informal)**.** *There exists a quantum oracle relative to which PRSGs exist, but PD-OWFs do not.*

As mentioned earlier, an immediate consequence of the above is that relative to a quantum oracle, PRSGs exist but short-PRSGs do not, thus we cannot shrink the output length of a PRSG in a black-box way.

For a pictorial representation of the relations between the primitives proven by the aforementioned theorem see figure 7.1.

Relation to the literature. As mentioned earlier, the study of pseudo-determinism was initiated by Ananth, Lin and Yuen [ALY24]. In their work, they differentiate between fractions of the key space where the algorithm does not behave deterministically; if this fraction is inverse polynomial in size, they denote this by inverse polynomial-error pseudo-determinism (PD-PRGs$^{\mathsf{poly}}$), and they analogously define negligible-error pseudo-determinism when this fraction is negligible in size (PD-PRGs$^{\mathsf{negl}}$). They showed that short-PRSGs can be used to construct PD-PRGs$^{\mathsf{poly}}$, and that PD-PRSGs$^{\mathsf{negl}}$ are enough to build PRSGs, as these can be used to generate pseudorandom phases pseudo-deterministically. However, it is still unclear how to build negligible-error PD-PRGs from short-PRSGs, or if inverse polynomial-error PD-PRGs are enough to build PRSGs[1]. In other words, there seems to be a gap between the cryptography we can build from inverse polynomial or negligible pseudo-determinism. For some applications, the polynomial error is not a big problem, as this error can be dealt with by repetition to construct commitments and encryption [ALY24], or using a recognizable abort to construct signatures [BBO+24].

In this chapter we show that even inverse polynomial error is not attainable from PRSGs in a black-box way, and hence, the output length of PRSGs cannot be shrunk. Our result was proven independently by Chung, Goldin and Gray [CGG24]. In their paper, they study One Way Puzzles (OWPs), and show a black-box separation between general OWPs and efficiently verifiable OWPs. This separation also implies a separation between short-PRSGs and PRSGs, because PRSGs imply OWPs and short-PRSGs imply efficiently verifiable OWPs. A concurrent work by Coladangelo and Mutreja [CM24] separating quantum digital signatures from PRSGs also provides a separation between short-PRSGs and PRSGs, since short-PRSGs imply quantum digital signatures as proven by Barhoush et al. [BBO+24].

7.2 Preliminaries

Although the general quantum-information notation can be found in Chapter 2, here we will introduce some concepts and prior results that we need only in this chapter. In particular, we discuss general properties of Haar-random states and give the definitions of quantum pseudorandom primitives we will be working with in the rest of the chapter. Moreover, we include some relevant results in complexity theory.

7.2.1 Notation

We use ε to denote the empty string. We use $x \prec y$ to denote the fact that x is a prefix of y, i.e. there exists x' such that $y = x||x'$. We use $x \not\prec y$ to denote the fact that x is not a prefix of y.

7.2.2 Definitions

We include here the relevant quantum pseudorandom primitives for this chapter. Recall the quantum pseudorandom state generators (PRSG) from Defi-

[1] The construction of PRSGs from PRGs does not work for PD-PRGs$^{\mathsf{poly}}$, as the polynomial error will induce too much error in the resulting state to be a PRS.

nition 6.2.1, with $\ell(\lambda) = \omega(\log \lambda)$-qubit outputs. We denote by *short*-PRSG a PRSG with $\ell(\lambda) := \Omega(\log \lambda)$-qubit output generator for some constant $c > 1$.

We will also work with the following pseudo-deterministic primitive, a variant of PD-PRGs from Definition 6.2.4.

Definition 7.2.1 (Quantum Pseudo-deterministic One-Way Functions[2])**.** *Let $\lambda \in \mathbb{N}$. Let m and ℓ be two polynomials. A* QPT *algorithm $F\colon \{0,1\}^\lambda \times \{0,1\}^{m(\lambda)} \to \{0,1\}^{\ell(\lambda)}$, often simplified as the family $F_\lambda := F(1^\lambda, \cdot)$ for $\lambda \in \mathbb{N}$, is a* quantum pseudo-deterministic one-way function *(PD-OWF) if the following conditions hold:*

i. ***Pseudo-determinism:*** *There exists a constant $c > 0$ and function $\mu(\lambda) := O(\lambda^{-c})$ such that for all $\lambda \in \mathbb{N}$, there exists a set $\mathcal{K}_\lambda \subset \{0,1\}^{m(\lambda)}$ with*

(a) $\Pr_{x \leftarrow \{0,1\}^{m(\lambda)}}(x \in \mathcal{K}_\lambda) \geq 1 - \mu(\lambda)$.

(b) For any $x \in \mathcal{K}_\lambda$, it holds that

$$\max_{y \in \{0,1\}^{\ell(\lambda)}} \Pr(y = F_\lambda(x)) \geq 1 - \mathsf{negl}(\lambda)\,, \tag{7.1}$$

where the probability is over the randomness of F_λ.

ii. ***Security:*** *For every* QPT *inverter $\mathcal{A}$:*

$$\Pr_{x \leftarrow \{0,1\}^{m(\lambda)}}(F_\lambda(\mathcal{A}(F_\lambda(x))) = F_\lambda(x)) \leq \mathsf{negl}(\lambda)\,, \tag{7.2}$$

where the probability is over the randomness of F_λ and $\mathcal{A}$.

Note that the pseudo-determinism factor in the above definition comes from the size of the *good* key space $\mu(\lambda)$, which is an inverse-polynomial in the security parameter λ. This means that for a non-negligible number of elements in the key space, the OWF could behave arbitrarily. We could also define a negligible variant by requesting $\mu(\lambda)$ to be a negligible function in λ. If the size of the key space is relevant, we can make it explicit by denoting $\mathrm{PD}-\mathrm{OWF}^{\mathsf{poly}}$ for inverse-polynomial $\mu(\lambda)$, or $\mathrm{PD}-\mathrm{OWF}^{\mathsf{negl}}$ for negligible $\mu(\lambda)$. And the same holds for $\mathrm{PD}-\mathrm{PRG}^{\mathsf{poly}}$ and $\mathrm{PD}-\mathrm{PRG}^{\mathsf{negl}}$.

We can build PD-OWFs from short-PRSs.

[2] In [BBO+24] they actually define Quantum Pseudo-deterministic One-Way *Hash* Functions (PD-OWHF). We omit the *hash* property here for simplicity, but since the security properties of both functionalities are equivalent our proof also trivially works for PD-OWHF.

Theorem 7.2.2 (Adapted from [BBO+24, Theorem 6][3]). *Assuming the existence of $(c\log\lambda)$-PRSs with $c > 12$, there exists a PD-OWF $F_\lambda : \{0,1\}^{\ell(\lambda)} \to \{0,1\}^{\ell(\lambda)}$ with input/output length $\ell(\lambda) = \lambda^{c/6}$ and key space of size $\mu(\lambda) = O(\lambda^{-c/12+1})$.*

7.2.3 Complexity theory

Definition 7.2.3. *A promise problem Π is in* PromiseBQP *if there exists a randomized polynomial-time quantum algorithm $\mathcal{A}$ such that, for all inputs $x \in \Pi_{\mathrm{Y}} \cup \Pi_{\mathrm{N}}$,*

$$\Pr(\mathcal{A}(x) = \chi_\Pi(x)) \geq \frac{2}{3}\,.$$

Definition 7.2.4. *A promise problem Π is in* PromiseQMA *if there exists a polynomial-time quantum algorithm* V *called the* verifier *and a polynomial p such that:*

i. ***Completeness.*** *If $x \in \Pi_{\mathrm{Y}}$, then there exists a quantum state $|\psi\rangle$ on $p(|x|)$ qubits (called a* witness*) such that*

$$\Pr(\mathsf{V}(x, |\psi\rangle) = 1) \geq \frac{2}{3}\,. \tag{7.3}$$

ii. ***Soundness.*** *If $x \in \Pi_{\mathrm{N}}$, then for every quantum state $|\psi\rangle$ on $p(|x|)$ qubits,*

$$\Pr(\mathsf{V}(x, |\psi\rangle) = 1) \leq \frac{1}{3}\,. \tag{7.4}$$

We talk about *promise problems* to emphasize that there might be instances that are neither YES nor NO, i.e. $\{0,1\}^* \setminus (\Pi_{\mathrm{Y}} \cup \Pi_{\mathrm{N}}) \neq \emptyset$. We do however drop the adjective when this is not the case, and talk simply about BQP and QMA problems. We say that the problem Π is in QCMA if the witness $|\psi\rangle$ is a classical state.

Theorem 7.2.5 ([Kre21]). *Let $\mathcal{U} = \{U_n\}_{n\in\mathbb{N}}$ be a quantum oracle where each U_n is chosen randomly from μ_{2^n}. Let $\mathcal{P}$ be an arbitrary* PSPACE*-complete language. There exists a quantum oracle $\mathcal{O} := (\mathcal{U}, \mathcal{P})$, such that:*

[3] Here we also use the PD-OWF variant of their theorem originally for PD-OWHF. This choice affects the parameters of the domain and range in the theorem statement because constructing a PD-OWHF requires more steps than constructing a PD-OWF (we only need the first step of their proof). However, note that changing the domain/range of the function to some different polynomials in λ would still make the proof go through by changing some parameters in the proof.

i. $\mathsf{PromiseBQP}^{\mathcal{O}} = \mathsf{PromiseQMA}^{\mathcal{O}}$.

ii. *λ-PRSGs exist relative to $\mathcal{O}$.*

7.3 No shrinking of PRSGs

In this section we will prove that the output length of pseudorandom state generators (PRSGs) cannot be shrunk in a black-box way. On the one hand, this intuitively might make sense, because if we naively try to throw away half of a maximally-entangled pure state (as is the case of Haar-random states), we might end up with a maximally-mixed state. On the other hand, there might have been cleverer solutions.

With our oracle separation we are saying that having *small* random looking states is possibly harder than having *larger* random looking states, or at least that we cannot obtain smaller random-looking states from larger ones.

Theorem 7.3.1. *There exists a quantum oracle $\mathcal{O}$ such that relative to it, λ-PRSGs exist, but $(c\log\lambda)$-PRSGs with $c > 12$ do not exist.*

Our proof goes through pseudo-determinism. The oracle necessary for the separation is actually the same oracle that Kretschmer used to separate PRSGs and OWFs in Theorem 7.2.5, as it turns out that relative to this oracle $\mathsf{PromiseBQP} = \mathsf{PromiseQMA}$, but then pseudo-deterministic one-way functions (PD-OWFs) do not exist either. In combination to Barhoush, Behera, Ozer, Salvail and Sattath's Theorem 7.2.2 result that short-PRSGs are enough to build PD-OWFs, we obtain the theorem.

According to the above considerations, the main theorem follows directly from the following one.

Theorem 7.3.2. *If PD-OWFs exist relative to a quantum oracle $\mathcal{O}$, then necessarily $\mathsf{PromiseBQP}^{\mathcal{O}} \neq \mathsf{PromiseQMA}^{\mathcal{O}}$.*

Proof. Let $\lambda \in \mathbb{N}$ and $F_\lambda \colon \{0,1\}^* \to \{0,1\}^{\ell(\lambda)}$ be a PD-OWF relative to the oracle $\mathcal{O}$. Let us define a promise problem Π, where YES instances have a pre-image with respect to F_λ but NO instances do not. Formally,

$$\Pi_{\mathrm{Y}} := \Big\{(1^{\ell(\lambda)}, x', y) \in \{1^{\ell(\lambda)}\} \times \{0,1\}^* \times \{0,1\}^{\ell(\lambda)} \Big| \\ \exists x \in \{0,1\}^*, x' \prec x, \Pr(y = F_\lambda(x)) \geq 1 - \mathsf{negl}(\lambda)\Big\}\,, \quad (7.5)$$

$$\Pi_{\mathrm{N}} := \left\{ (1^{\ell(\lambda)}, x', y) \in \{1^{\ell(\lambda)}\} \times \{0,1\}^* \times \{0,1\}^{\ell(\lambda)} \middle| \right.$$
$$\left. \forall x \in \{0,1\}^*, x' \not\prec x \text{ or } \Pr(y = F_\lambda(x)) \leq 1 - \frac{1}{\lambda} \right\}. \quad (7.6)$$

Note that in the definition of Π_{Y}, we have that $\Pr(y \neq F_\lambda(x))$ is a negligible function, and we will use this property in the rest of the section. We claim that

Claim 7.3.1. *The promise problem* Π *is in* $\mathsf{PromiseQMA}^{\mathcal{O}}$.

Claim 7.3.2. *The promise problem* Π *is not in* $\mathsf{PromiseBQP}^{\mathcal{O}}$.

Therefore, there must be a separation between both complexity classes. □

We now prove the two claims from the proposition. We start by showing that the language defined in Theorem 7.3.2 is in $\mathsf{PromiseQMA}^{\mathcal{O}}$, that is, we will construct an algorithm (verifier) that given an element $x \in \Pi_{\mathrm{Y}} \cup \Pi_{\mathrm{N}}$ and a (quantum) proof, can distinguish if the element is a YES or NO instance.

Proof of Claim 7.3.1. We define a quantum polynomial-time algorithm $\mathcal{A}^{\mathcal{O}}$ that given an element of the domain $(1^{\ell(\lambda)}, x', y) \in \{1^{\ell(\lambda)}\} \times \{0,1\}^* \times \{0,1\}^{\ell(\lambda)}$ and a classical proof $x \in \{0,1\}^{\ell(\lambda)}$, will check if the proof x is indeed a preimage of the PD-OWF by checking if it coincides with the input y.

Algorithm 22 $\mathcal{A}^{\mathcal{O}}((1^{\ell(\lambda)}, x', y), x)$

```
if $x' \not\prec x$ then return 0
end if
for $1 \leq i \leq 2\lambda$ do
    if $F_\lambda(x) \neq y$ then return 0
    end if
end for
return 1
```

Note that the algorithm runs in polynomial time trivially because computing F_λ is done efficiently relative to $\mathcal{O}$ by definition and we make 2λ calls to it. We now prove that the algorithm distinguishes between the YES/NO instances.

Case 1: Let $(1^{\ell(\lambda)}, x', y) \in \Pi_{\mathrm{Y}}$. Then by definition there exists a proof $x \in \{0,1\}^*$ such that

$$x' \prec x \text{ and } \Pr(y = F_\lambda(x)) \geq 1 - \mathsf{negl}(\lambda)\,. \tag{7.7}$$

Then the proof x will be an element of the input $((1^{\ell(\lambda)}, x', y), x)$ for which the algorithm $\mathcal{A}$ will output 1 with high probability because

$$\begin{aligned}\Pr\Big(\mathcal{A}^{\mathcal{O}}((1^{\ell(\lambda)}, x', y), x) = 1\Big) &= \Pr(\forall 1 \leq i \leq 2\lambda,\, y = F_\lambda(x)) && (7.8)\\ &\geq (1 - \mathsf{negl}(\lambda))^{2\lambda} \geq 2/3\,, && (7.9)\end{aligned}$$

which holds whenever λ is big enough. Indeed, recall that $\mathsf{negl}(\lambda) \leq 1/\lambda^c$ for all $c > 1$ and λ big enough, thus in particular $\mathsf{negl}(\lambda) \leq 1/\lambda^2$ for λ big enough, hence

$$(1 - \mathsf{negl}(\lambda))^{2\lambda} \geq \left(1 - \frac{1}{\lambda^2}\right)^{2\lambda} \geq \frac{2}{3}\,, \tag{7.10}$$

whenever λ is big enough. For the last inequality, we need at least $\lambda \geq 6$, where we used that we have an increasing function in λ.

Case 2: Let $(1^{\ell(\lambda)}, x', y) \in \Pi_{\mathrm{N}}$. Then by definition for every potential proof $x \in \{0,1\}^*$ we have that either

$$x' \not\prec x \quad \text{or} \quad \Pr(y = F_\lambda(x)) \leq 1 - \frac{1}{\lambda}\,. \tag{7.11}$$

Then for every possible input $((1^{\ell(\lambda)}, x', y), x)$ the algorithm will output 0 with high probability because

$$\begin{aligned}\Pr\Big(\mathcal{A}((1^{\ell(\lambda)}, x', y), x) = 1\Big) &= \Pr\big(x' \prec x \wedge \forall 1 \leq i \leq 2\lambda,\, y = F_\lambda(x)\big) && (7.12)\\ &\leq \Pr(\forall 1 \leq i \leq 2\lambda,\, y = F_\lambda(x)) && (7.13)\\ &\leq \left(1 - \frac{1}{\lambda}\right)^{2\lambda} \leq e^{-2} \leq 1/3\,. && (7.14)\end{aligned}$$

□

Proof of Claim 7.3.2. We will prove this by contradiction. Let us assume instead that Π is in $\mathsf{PromiseBQP}^{\mathcal{O}}$, that is, there exists a BQP algorithm $\mathcal{A}^{\mathcal{O}}$ such that for all $(1^{\ell(\lambda)}, x', y) \in \Pi_{\mathrm{Y}} \cup \Pi_{\mathrm{N}}$:

$$\Pr\Big(\mathcal{A}((1^{\ell(\lambda)}, x', y)) = \chi_\Pi((1^{\ell(\lambda)}, x', y))\Big) \geq \frac{2}{3}\,. \tag{7.15}$$

Without loss of generality, we can assume that the algorithm $\mathcal{A}^{\mathcal{O}}$ has completeness $1-\frac{1}{\ell(\lambda)}$ and soundness $\frac{1}{\ell(\lambda)}$. We will now show how we can construct a QPT algorithm $\mathcal{A}'^{\mathcal{O}}$ that finds a pre-image of every F_λ with high probability when it exists, by querying the original BQP algorithm $\mathcal{A}^{\mathcal{O}}$ at most $\ell(\lambda)+1$ times.

Algorithm 23 $\mathcal{A}'^{\mathcal{O}}(1^{\ell(\lambda)}, y)$

```
$b \leftarrow \mathcal{A}^{\mathcal{O}}(1^{\ell(\lambda)}, \varepsilon, y)$
if $b = 0$ then return $\perp$
end if
$x_0 \leftarrow \varepsilon$
for $1 \leq i \leq \ell(\lambda)$ do
    $b \leftarrow \mathcal{A}^{\mathcal{O}}(1^{\ell(\lambda)}, x_0||0, y)$
    if $b = 1$ then
        $x_0 = x_0||0$
    else
        $x_0 = x_0||1$
    end if
end for
return $x_0$
```

Indeed if $(1^{\ell(\lambda)}, \varepsilon, y) \in \Pi_{\mathrm{Y}}$, then the probability that the algorithm $\mathcal{A}'^{\mathcal{O}}$ outputs a correct pre-image is very high

$$\Pr\Big(y = \operatorname{argmax}_{y\in\{0,1\}^{\ell(\lambda)}} \Pr(y = F_\lambda(x)) \,\Big|\, x \leftarrow \mathcal{A}'^{\mathcal{O}}\big(1^{\ell(\lambda)}, y\big)\Big) \geq \left(1 - \frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1}. \tag{7.16}$$

However, this raises a contradiction with the security of the PD-OWF from the assumption of Definition 7.2.1,

$$\Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\Big(F_\lambda(\mathcal{A}'^{\mathcal{O}}(1^{\ell(\lambda)}, F_\lambda(x))) = F_\lambda(x)\Big) \tag{7.17}$$

$$= \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\Big(F_\lambda(\mathcal{A}'^{\mathcal{O}}(1^{\ell(\lambda)}, F_\lambda(x))) = F_\lambda(x) \,\Big|\, x \in \mathcal{K}_\lambda\Big) \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}(x \in \mathcal{K}_\lambda) \tag{7.18}$$

$$+ \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\Big(F_\lambda(\mathcal{A}'^{\mathcal{O}}(1^{\ell(\lambda)}, F_\lambda(x))) = F_\lambda(x) \,\Big|\, x \notin \mathcal{K}_\lambda\Big) \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}(x \notin \mathcal{K}_\lambda) \tag{7.19}$$

$$\geq \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left(F_\lambda(\mathcal{A}'^{\mathcal{O}}(1^{\ell(\lambda)},F_\lambda(x)))=F_\lambda(x)\,\middle|\, x\in\mathcal{K}_\lambda\right)\Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}(x\in\mathcal{K}_\lambda) \tag{7.20}$$

$$\geq \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left(F_\lambda(\mathcal{A}'^{\mathcal{O}}(1^{\ell(\lambda)},F_\lambda(x)))=F_\lambda(x)\,\middle|\, x\in\mathcal{K}_\lambda\right)(1-\mu(\lambda))\,, \tag{7.21}$$

where the first equality comes from the law of total probability and the second inequality comes from the property of $\mathcal{K}_\lambda$. We can rewrite the probability of the last equation as:

$$\Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left(y_3=y_1\middle| \begin{array}{r} y_1,y_2\leftarrow F_\lambda(x), x_1\leftarrow\mathcal{A}'^{\mathcal{O}}(1^{\ell(\lambda)},y_2),\\ y_3\leftarrow F_\lambda(x_1), x\in\mathcal{K}_\lambda\end{array}\right) \tag{7.22}$$

$$\geq \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left(\left(y_3=y_2=y_1=\mathrm{argmax}_{y\in\{0,1\}^{\ell(\lambda)}}\Pr[y=F_\lambda(x)]\right)\wedge(x_1=x)\right.\left.\middle|\begin{array}{c} y_1,y_2,y_3\leftarrow F_\lambda(x)\\ x_1\leftarrow\mathcal{A}'^{\mathcal{O}}(1^{\ell(\lambda)},y_2), x\in\mathcal{K}_\lambda\end{array}\right) \tag{7.23}$$

$$= \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left(F_\lambda(x)=\mathrm{argmax}_{y\in\{0,1\}^{\ell(\lambda)}}\Pr[y=F_\lambda(x)]\,\middle|\, x\in\mathcal{K}_\lambda\right)^3 \cdot \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left(\mathcal{A}'^{\mathcal{O}}\left(1^{\ell(\lambda)},\mathrm{argmax}_{y\in\{0,1\}^{\ell(\lambda)}}\Pr[y=F_\lambda(x)]\right)=x\,\middle|\, x\in\mathcal{K}_\lambda\right) \tag{7.24}$$

$$\geq (1-\mathsf{negl}(\lambda))^3\left(1-\frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1}, \tag{7.25}$$

where the first inequality comes from the law of total probability, and the last inequality from the definition of a PD-OWF and eq. (7.16). This gives that:

$$\Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left(F_\lambda(\mathcal{A}'^{\mathcal{O}}(1^{\ell(\lambda)},F_\lambda(x)))=F_\lambda(x)\right) \tag{7.26}$$

$$\geq (1-\mathsf{negl}(\lambda))^3\left(1-\frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1}(1-\mu(\lambda)) \tag{7.27}$$

$$\geq (1-O(\mu(\lambda)))\left(1-\frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1}. \tag{7.28}$$

Note that this bound is not negligible since

$$\left(1-\frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1}\geq\frac{1}{10}\,, \tag{7.29}$$

whenever $\lambda \geq 2$, which contradicts eq. (7.2). □

7.4 Conclusion

In this chapter we discussed quantum pseudorandomness, and proved a property that makes it different from classical pseudorandomness. Namely, the output length of the primitives is very relevant, as shown by the impossibility of shrinking the output length of quantum pseudorandom states.

Finding a unifying framework for quantum pseudorandomness will not be as straightforward as for classical pseudorandomness, where one-way functions are a central resource. However, it is still not clear what is the underlying reason for their difference, and thus makes quantum pseudorandomness an object of fascination. We include here some open questions related to the work in this chapter.

- Our *no-shrinking* result is relative to a quantum oracle, thus a *classical* oracle separation of short-PRSGs from PRSGs is still unclear. The classical oracle from [KQST23] gives $\mathsf{P} = \mathsf{NP}$, but this is not enough for PD-OWFs, as we need promise problems in our proof.

- Here we focus on the extreme case of short-PRSGs, but it is still open if we can discard a *linear* amount of qubits while maintaining security. That is, it could be that given a λ-PRS we could, for example, construct a $\lambda/2$-PRS.

- Oracle separations only provide black-box separations, thus it is still interesting to ask if for general PRSG constructions there are ways of shrinking the output length. For example, there exist ways of lengthening the output length of pseudorandom phase PRSGs [LV24].

- We can also wonder if the opposite of no-shrinking is true, that is, if we can lengthen the output size of PRSGs.

Chapter 8

Constructing Quantum Pseudorandom primitives

This chapter is based on the article *Cryptography from Lossy Reductions: Towards OWFs from ETH, and Beyond* [FGMR25], and is joint work with Alex B. Grilo, Pouria Fallahpour and Mahshid Riahinia.

The introduction has been rewritten to put the focus on quantum primitives, as we have omitted all the exclusively classical results from the original work. In particular, we have merged [FGMR25, Sections 6 and 7] and omitted the specific construction of classical one-way functions in [FGMR25, Sections 6]. From the original section [FGMR25, Section 8] we have omitted the results about classical reductions, and we have omitted [FGMR25, Section 9] as we do not discuss particular languages nor parameters for our reductions. A conclusion section has been added and the remaining sections have been edited for style and typographical consistency.

In this chapter we discuss the meta-complexity of quantum cryptography, that is, hardness requirements for quantum cryptographic primitives to exist. In Section 8.3 we introduce quantum lossy mappings and prove an extended *quantum disguising distribution* lemma. In Section 8.4 we introduce *mildly-lossy* problems, that is, promise problems that admit quantum lossy f-distinguisher reductions, for non-constant permutation-invariant functions f. In Section 8.5 we show that mildly-lossy problems admit Karp reductions to a problem in QSZK. We use the previous result to prove in Section 8.6 that we can build quantum cryptography from lossy problems; namely, efficiently samplable statistically far but computationally indistinguishable pairs of states (EFIs) and one-way state generators (OWSGs). Finally, in Section 8.7 we show that worst-case to average-case reductions are lossy.

8.1 Introduction

Efficiently samplable, statistically far but computationally indistinguishable pairs of states (EFIs) are a strong candidate for being the minimal assumption for *quantum* cryptography; the existence of EFIs implies quantum bit commitment, oblivious transfer, and multi-party computation [MY22, BCQ23]. Moreover, they are implied by all the other quantum primitives introduced in the last years such as quantum pseudorandom states or one-way state generators (OWSGs), classical-input pure-outcome quantum maps that are difficult to invert. While EFI pairs posit the existence of a problem that is hard to *distinguish*, OWSGs assume the existence of a hard *search* problem that admits an efficient verification procedure.

Given the interest on basing cryptography on these new quantum primitives, one would like to know if they *actually* exist. Ideally, one would like to build quantum cryptography unconditionally or under minimal complexity theoretical assumptions. Given that we know that the existence of classically hard-to-solve but easy-to-verify problems given a witness (i.e. $\mathsf{P} \neq \mathsf{NP}$) is not a necessary prerequisite for quantum cryptography [Kre21], in this thesis we ask the following question.

What are the minimal complexity theoretical assumptions for the existence of quantum cryptography?

Our contributions. Trying to answer the above question, we initiate the study of cryptographic implications of quantum *lossy f-distinguisher* reductions. A quantum reduction $R\colon \{0,1\}^* \to \mathsf{MS}_*$ is said to be $\ell(\lambda)$-lossy for a class of distributions $\{X_\lambda\}_{\lambda\in\mathbb{N}}$ when $I_q(X_\lambda; R(X_\lambda)) \leq \ell(\lambda)$, where I_q is the quantum mutual information. Moreover, such a reduction is said to be a pure-outcome reduction if for every instance x the outcome $R(x)$ is a pure quantum state. We define quantum f-*distinguisher* reductions as follows: let Π be a promise problem, and let χ_Π be the characteristic function of Π, that is, for an input x, $\chi_\Pi(x) = 1$ if x is a YES instance of Π, and 0 otherwise. For a function $f : \{0,1\}^m \to \{0,1\}$, a quantum f-distinguisher reduction R for Π is such that on input m instances of Π, there exists an unbounded algorithm $\mathcal{D}$ that can distinguish between $R(x_1,\ldots,x_m)$ and $R(x'_1,\ldots,x'_m)$, also given one of $\{x_i\}_i$'s at random, when $f(\chi_\Pi(x_1),\ldots,\chi_\Pi(x_m)) \neq f(\chi_\Pi(x'_1),\ldots,\chi_\Pi(x'_m))$.

Our main technical contribution is generalizing Drucker's [Dru12] *disguising distribution lemma*, informally stating that not much information about x can be recovered from $R(x)$. Our results hold with respect to quantum lossy f-distinguisher reductions R for any non-constant permutation-invariant function f, say Or, And, Maj, Parity, Mod$_k$ or Threshold$_k$. We call problems Π that admit such reductions R *mildly-lossy problems.*

We then apply the extended disguising lemma to build quantum cryptographic primitives from complexity-theoretic assumptions. In particular, we prove the following theorem:

Theorem 8.1.1 (informal)**.** *Let* Π *be a mildly-lossy problem:*

i. If the reduction is $m\alpha$*-lossy for a small constant* α*, then we construct a non-uniform polynomial-time quantum algorithm* EFI *that is a EFI scheme.*

ii. If the reduction has pure-outcome and is $m\alpha$*-lossy for a* larger α*, then we construct a tuple of non-uniform quantum algorithm* $\mathsf{G}_\Pi :=$ $(\mathsf{StateGen}, \mathsf{Ver})$*, where* $\mathsf{StateGen}$ *runs in polynomial-time, that is a one-way state generator.*

When the outcome of the reduction is classical, the above statements also hold for one-way functions.

Finally, we give evidence for the existence of mildly-lossy problems by analysing the mild-lossiness of special well-known reductions, worst-case to average-case Karp reductions, i.e. polynomial-time many-to-one reductions. In particular, we relate mild-lossiness to the distance between the output distribution of the worst-case to average-case reduction and the average-case distribution.

Theorem 8.1.2 (informal)**.** *Let* Π *be a promise problem that admits a worst-case to average-case Karp reduction:*

i. If the minimum eigenvalue of the output of a quantum reduction is uniformly bounded from below by a constant β*, then* Π *is a mildly-lossy problem, where the loss is linear in* β*.*

ii. If instead the dimension of the image of the reduction is upper bounded by a constant d*, then* Π *is a mildly-lossy problem, where the loss is logarithmic in* d*.*

The two conditions come from two different reverse Pinsker-like inequalities for quantum mutual information.

Relation to the literature. In contrast to the sea of cryptographic primitives introduced, few works have presented evidence to support their existence. Chailloux, Kerenidis and Rosgen [CKR11] initiated the task of basing the existence of quantum cryptographic protocols on worst-case complexity assumptions. They proved that if QSZK (the class of all languages that have an interactive proof where a quantum verifier learns almost nothing beyond the membership of the instance in the language) is worst-case hard, then auxiliary-input EFIs exist. Very recently, Morimae, Nehoran and Yamakawa [MNY24] eliminated the hardness assumption by showing that auxiliary-input EFIs exist unconditionally given quantum advice. Nevertheless, one would like to build cryptographic primitives without auxiliary inputs. Towards this direction Bostanci, Haferkamp, Hangleiter and Poremba [BHHP24] and concurrently Hiroka and Morimae [HM24] based the existence of EFIs on the hardness of two concrete problems: the *Hamiltonian Phase State problem* and GapK problem, respectively. This chapter not only bases the existence of EFIs, without auxiliary input, on properties of quantum reductions, but also gives conditions to extend this to pure-outcome one-way state generators (a possibly stronger primitive than EFIs, as discussed in Chapter 6).

Chapter overview. Since the chapter is technically involved we include here an overview of its structure. In Section 8.3 we prove an extended version of Drucker's [Dru12] disguising distribution lemma; we extend it from quantum compressing mappings to quantum *splitting-lossy*. In Section 8.4 we introduce f-distinguisher reductions for promise problems Π; a mapping that sends m instances $x_1, \ldots, x_m$ of size λ to a quantum state ρ, such that no distinguisher can decide $f(\xi_\Pi(x_1), \ldots, \xi_\Pi(x_m))$ for some Boolean function f. In this chapter we also define *mildly-lossy problems*, which are informally promise problems that admit a f-distinguisher splitting lossy reduction. In Section 8.5, we give a range of parameters for a mildly-lossy problem to reduce to a problem in QSZK (or SZK when the reduction is classical), this is done using the completeness of the quantum state distinguishability problem in QSZK, which gives a pair of circuits with the promise that either their evaluations are close or far in trace distance. In Section 8.6, we leverage the structure of the reduction to QSZK, to construct EFI schemes and OWSGs from mildly lossy problems. Finally, in Section 8.7, we give arguments to support the abstraction to mildly-lossy problems by giving conditions for worst-case to average-case reductions to be mildly lossy.

8.2 Preliminaries

Although the general quantum-information notation can be found in Chapter 2, here we will introduce some concepts and prior results that we need only in this chapter. In this chapter, we always consider non-uniform algorithms. All classical algorithms are quantum algorithms, therefore, we mostly use the quantum formalism for generalization and simplification. When the distinction is necessary, we explicitly mention it in the beginning of a section or inside a statement, and clearly distinguish between classical and quantum settings.

8.2.1 Notation

The set of all permutations over $[n]$ is $\mathfrak{S}_n$. We abuse notation and use the same symbol to refer to the uniform distribution over all permutations of $[n]$. For any finite set S, we let $\mathcal{U}_S$ denote the uniform distribution over S. A distribution is called s-uniform if it is sampled uniformly from a multiset of at most s elements.
A Boolean function $f : \{0,1\}^m \to \{0,1\}$ is called non-constant if it is not always 0 nor always 1. Given two languages $\mathcal{L}, \mathcal{L}' \subseteq \{0,1\}^*$ and a computable function $f \colon \mathcal{L} \to \mathcal{L}'$, f is said to be a *many-to-one reduction* if for every $x \in \{0,1\}^*$, $x \in \mathcal{L}$ if and only if $f(x) \in \mathcal{L}'$.

8.2.2 Games

For two non-zero natural numbers $a, b \in \mathbb{N}^+$, a two-player, simultaneous-move, zero-sum game is specified by a matrix $\mathbf{M} \in \mathbb{R}^{a \times b}$. Player 1 chooses a row index $i \in [a]$ and Player 2 chooses a column index $j \in [b]$, and Player 2 receives the payoff $\mathbf{M}_{ij}$ from Player 1. The goal of Player 1 is minimizing the expected payoff, while Player 2 opts to maximize it. The row and column indices are called the pure strategies of Player 1 and Player 2, respectively. The mixed strategies are distributions over possible choices of indices. A mixed strategy is s-uniform if it is sampled uniformly from a multiset of at most s pure strategies.

Lemma 8.2.1 ([vN28]). *Let $\mathcal{P}$ and $\mathcal{Q}$ be two mixed strategies for Player 1 and 2, respectively. It holds that* $\min_{\mathcal{P}} \max_j \mathbb{E}_{i \sim \mathcal{P}}[\mathbf{M}_{ij}] = \max_{\mathcal{Q}} \min_i \mathbb{E}_{j \sim \mathcal{Q}}[\mathbf{M}_{ij}]$.

The value of the game, which we denote by $\omega(\mathbf{M})$, is the optimal expected value guaranteed by the above lemma. The following lemma shows that each player has nearly-optimal s-uniform strategy when s is chosen to be the logarithm of the number of pure strategies of the opponent.

Lemma 8.2.2 ([LY94, Theorem 2])**.** *For any real $\varepsilon > 0$, $a, b \in \mathbb{N}^+$, $M \in \mathbb{R}^{a\times b}$, and any integer $s \geq \ln(b)/(2\varepsilon^2)$, it holds that*

$$\min_{\mathcal{P}\in\mathfrak{P}_s} \max_{j} \mathbb{E}_{i\sim\mathcal{P}}[\mathbf{M}_{ij}] \leq \omega(\mathbf{M}) + \varepsilon(\mathbf{M}_{max} - \mathbf{M}_{min}) \ , \tag{8.1}$$

where $\mathfrak{P}_s$ denotes the set of all s-uniform strategies for Player 1, $\mathbf{M}_{min} := \min_{i,j} M_{ij}$ and $\mathbf{M}_{max} := \max_{i,j} M_{ij}$. A similar statement holds for Player 2, namely,

$$\max_{\mathcal{Q}\in\mathfrak{Q}_s} \min_{i} \mathbb{E}_{j\sim\mathcal{Q}}[\mathbf{M}_{ij}] \geq \omega(\mathbf{M}) - \varepsilon(\mathbf{M}_{max} - \mathbf{M}_{min}) \ , \tag{8.2}$$

where $\mathfrak{Q}_s$ denotes the set of all s-uniform strategies for Player 2.

8.2.3 Information theory

Classical information. Given two probability distributions X and Y over Σ, their statistical distance, also called total variation distance, is defined as

$$\Delta(X,Y) := \frac{1}{2}\sum_{x\in\Sigma} |\Pr(X = x) - \Pr(Y = x)| \, . \tag{8.3}$$

The Kullback-Leibler divergence or classical relative entropy of X with respect to Y is defined as

$$D_{KL}(X||Y) := \sum_{x\in\Sigma} \Pr(X = x) \log\left(\frac{\Pr(X = x)}{\Pr(Y = x)}\right) . \tag{8.4}$$

Quantum information. Let $R : \{0,1\}^n \to \mathsf{MS}_m$ be any quantum mapping and X a random variable supported over $\{0,1\}^n$. We let

$$\rho_{X,R(X)} := \sum_{x\in\{0,1\}^n} \Pr_X(x)|x\rangle\langle x| \otimes R(x) \ . \tag{8.5}$$

For a mixed state ρ, we let $S(\rho) := \mathrm{Tr}(\rho \log_2 \rho)$ denote the Von Neumann entropy of ρ. The quantum mutual information of two subsystems A and B is defined as

$$I_q(A;B)_\rho := S(\rho_A) + S(\rho_B) - S(\rho_{AB}) \ , \tag{8.6}$$

where $\rho_A = \mathrm{Tr}_B(\rho_{AB})$ and $\rho_B = \mathrm{Tr}_A(\rho_{AB})$. For the sake of simplicity, we sometimes drop the subscripts q and ρ in I_q. When working with quantum systems A, B, the notation $I(A; B)$ implicitly refers to $I_q(A; B)$.

For two quantum states ρ and σ, the quantum relative entropy of ρ with respect to σ is

$$D(\rho\|\sigma) := \begin{cases} \mathrm{Tr}(\rho(\log(\rho) - \log(\sigma))) & \text{if } \mathrm{Supp}(\rho) \subseteq \mathrm{Supp}(\sigma)\,, \\ \infty & \text{otherwise}\,. \end{cases} \tag{8.7}$$

Given a bipartite state ρ_{AB} with marginals ρ_A and ρ_B, the relative entropy can be written in terms of the mutual information as

$$D(\rho_{AB}\|\rho_A \otimes \rho_B) = I_q(A; B)_\rho\,. \tag{8.8}$$

For pairs of classical-quantum states of the form $\rho_{XB} := \sum_x p(x)|x\rangle\langle x|_X \otimes \rho_B^x$ and $\sigma_{XB} := \sum_x q(x)|x\rangle\langle x|_X \otimes \sigma_B^x$, the relative entropy takes the simpler form

$$D(\rho_{XB}\|\sigma_{XB}) = \sum_x p(x) D(\rho_B^x\|\sigma_B^x) + D_{KL}(p\|q)\,, \tag{8.9}$$

where D_{KL} is the classical Kullback-Leibler divergence.

Lemma 8.2.3 ([AE11, Theorem 1])**.** *Let ρ and σ be two quantum states, let the smallest eigenvalue of σ be uniformly bounded from below, i.e. there exists $\beta > 0$ such that $\lambda_{min}(\sigma) > \beta$. Then the relative entropy of ρ with respect to σ is bounded by*

$$D(\rho\|\sigma) \leq (\beta + T(\rho, \sigma)) \log\left(1 + \frac{T(\rho, \sigma)}{\beta}\right).$$

We let $S(\rho\|\sigma) := \mathrm{Tr}(\rho(\log(\rho) - \log(\sigma)))$ denote the relative entropy. We define the quantum conditional entropy of a two-system state ρ_{AB} as follows

$$S(A|B) := S(\rho_{AB}) - S(\rho_B)\,, \tag{8.10}$$

The quantum mutual information in terms of conditional quantum entropy is

$$I(A; B)_\rho = S(\rho_A) - S(A|B)_\rho = S(\rho_B) - S(B|A)_\rho\,. \tag{8.11}$$

Lemma 8.2.4. *Let ρ_{AB} be a quantum state in two subsystems A and B, with marginal states ρ_A and ρ_B. The following properties hold.*

i. The Von Neumann entropy is additive for tensor product states: $S(\rho_A \otimes \rho_B) = S(\rho_A) + S(\rho_B)$ *.*

ii. Conditioning does not increase entropy: $S(\rho_A) \geq |S(A|B)_\rho|$ *.*

iii. The quantum entropy of a system is bounded by the dimension: $S(\rho_A) \leq \dim(H_A)$ *.*

Lemma 8.2.5 (Alicki–Fannes–Winter Inequality [Wil13])**.** *Let* $\rho_{AB}, \omega_{AB} \in \mathcal{D}(H_A \otimes H_B)$*, then*

$$|S(A|B)_\rho - S(A|B)_\omega| \leq 2\mathrm{TD}(\rho, \sigma) \log \dim(H_A) + h(\mathrm{TD}(\rho, \sigma)) ,$$

where h *is the binary entropy function and* TD *refers to the trace distance.*

The following lemma states that if the outcome of a measurement is close to deterministic, then it must not alter much the state.

Lemma 8.2.6 (Gentle Measurement Lemma [Win99])**.** *Let* ρ *be a mixed state and* $\{\Pi, I - \Pi\}$ *a two-outcome POVM with* $\mathrm{Tr}(\Pi\rho) \geq 1 - \varepsilon$*, then* $\|\rho - \rho'\|_1 \leq \sqrt{\varepsilon}$*, where* $\rho' = \frac{\sqrt{\Pi}\rho\sqrt{\Pi}}{\mathrm{Tr}(\Pi\rho)}$*.*

Lemma 8.2.7 (Swap test)**.** *The swap test on input* (σ, ρ) *outputs* 1 *with probability*

$$\frac{1 + \mathrm{Tr}(\rho\sigma)}{2} ,$$

in which case we say that it passes the swap test.

Definition 8.2.8 (ℓ_1 distance for classical distributions and quantum states)**.** *We use the notation* $\|X - Y\|_1$ *to refer to (i) either statistical distance* $\Delta(X, Y)$ *when variables* X, Y *are classical distributions (ii) or trace distance* $\mathrm{TD}(X, Y)$ *when they are quantum states.*

8.2.4 Complexity theory

Search Problems. We recall the definition of a search problem, inspired by that of [BG94]. We define a *search* problem Π_{search} as a binary relation over $\{0,1\}^* \times \{0,1\}^*$. For any $(x, w) \in \Pi_{\mathsf{search}}$, we call x an *instance* and w a *witness*. For any $x \in \{0,1\}^*$, we define $\Pi_{\mathsf{search}}(x) = \{w \in \{0,1\}^* \colon (x, w) \in \Pi_{\mathsf{search}}\}$. We refer to the sets $\Pi_{\mathsf{search}_{|Y}} = \{x \in \{0,1\}^* \colon \Pi_{\mathsf{search}}(x) \neq \emptyset\}$, and $\Pi_{\mathsf{search}_{|N}} = \{0,1\}^* \backslash \Pi_{\mathsf{search}_{|Y}}$ as the set of YES and NO instances, respectively.

We say that an algorithm $\mathcal{A}$ solves Π_{search}, if for any $x \in \{0,1\}^*$ for which $\Pi_{\text{search}}(x) \neq \emptyset$, $\mathcal{A}$ returns some $w \in \Pi_{\text{search}}(x)$, and otherwise, outputs $\perp$.

We denote the decision language defined by Π_{search} as

$$\Pi = \{x \in \{0,1\}^* \colon \exists w \in \{0,1\}^*, (x,w) \in \Pi_{\text{search}}\} . \tag{8.12}$$

Each decision language Π can have multiple associated *search problems*, one for every relation Π_{search} that defines Π. Given $x \in \Pi$, the Π_{search}-search problem consists of finding $\omega \in \Pi_{\text{search}}(x)$.

Worst-case hardness. In this work, we consider fine-grained worst-case hardness, as introduced below.

Definition 8.2.9. *For a function $T : \mathbb{N} \to \mathbb{R}^+$, a promise problem Π is said to be $T(\lambda)$-hard, if for any non-uniform classical-advice algorithm $\mathcal{A}$ with runtime at most $T(\lambda)$ over λ-bit inputs, and any sufficiently large $\lambda \in \mathbb{N}$, there exists an input $x \in (\Pi_Y \cup \Pi_N) \cap \{0,1\}^\lambda$ such that $\Pr(\mathcal{A}(x) = \chi_\Pi(x)) < 2/3$.*

We will assume that the size of the advice is not larger than the runtime. By setting $T(\lambda) = \mathsf{poly}(\lambda)$, one recovers the regular definition of worst-case hardness.

Complexity class QSZK. We recall the quantum state distinguishability problem below. We refer to [Wat02] for more details.

Definition 8.2.10 (Quantum State Distinguishability)**.** *Let $\alpha, \beta \in [0,1]$ such that $\alpha < \beta$. Given two quantum circuits $\mathcal{C}_0$ and $\mathcal{C}_1$, let ρ_0 and ρ_1 be the (mixed) quantum states that they produce by running on all-zero states with the promise that either $\|\rho_0 - \rho_1\|_1 \geq \beta$ (corresponding to NO instances) or $\|\rho_0 - \rho_1\|_1 \leq \alpha$ (corresponding to YES instances). The $\mathrm{QSD}_{\alpha,\beta}$ problem is to decide which one is the case.*

The above problem enjoys a polarization property. The lemma below is adapted from [Wat02, SV03].

Lemma 8.2.11. *Let λ be a positive integer. Let $\alpha, \beta : \mathbb{N} \to [0,1]$, and $\theta : \mathbb{R} \to (1, +\infty)$ be functions of λ such that $\theta := \beta^2/\alpha$. There exists a deterministic classical algorithm $\mathsf{Polarize}$ that given a pair of (quantum) circuits (C_0, C_1) as well as a unary parameter 1^λ, outputs a pair of (quantum) circuits (P_0, P_1) such that*

$$\|C_0|\mathbf{0}\rangle - C_1|\mathbf{0}\rangle\|_1 \leq \alpha \Rightarrow \|P_0|\mathbf{0}\rangle - P_1|\mathbf{0}\rangle\|_1 \leq 2^{-\lambda} , \tag{8.13}$$

$$\|C_0|\mathbf{0}\rangle - C_1|\mathbf{0}\rangle\|_1 \geq \beta \Rightarrow \|P_0|\mathbf{0}\rangle - P_1|\mathbf{0}\rangle\|_1 \geq 1 - 2^{-\lambda} . \tag{8.14}$$

Moreover, the runtime and output size of $\mathsf{Polarize}$ *are of* $O(n\log(8\lambda)(|C_0|+|C_1|)/\log(\theta))$ *when* $\lambda \to +\infty$.

There are various equivalent definitions of the complexity class QSZK. The following definition suffices for our purposes.

Definition 8.2.12 (QSZK)**.** *The class* QSZK *consists of all promise problems that have many-to-one polynomial-time reductions to* $\mathrm{QSD}_{1/4,3/4}$.

All definitions and lemmas above can be restricted to classical algorithms. In this case, we let SZK denote the corresponding classical complexity class and SD denote the statistical difference problem (classical variant of QSD).

8.2.5 Cryptographic primitives

Recall the definitions of efficiently samplable statistically far but computationally indistinguishable quantum states (EFI) and one-way state generators (OWSG) from Chapter 6. We include here their non-uniform variants, as well as a refined description of their security which will help in understanding what type of primitives we are able to obtain in the later sections.

Definition 8.2.13 (Non-Uniform EFI)**.** *Let* $T : \mathbb{N} \to \mathbb{R}^+$ *and* $d, D : \mathbb{N} \to [0,1]$ *be functions. A non-uniform* (T, D, d)*-EFI scheme is a QPT algorithm* $\mathsf{EFI}_h(1^\lambda, b)$ *that given a classical* $\mathsf{poly}(\lambda)$*-size advice* h *and a bit* b*, outputs a quantum state* ρ_b*, such that for any sufficiently large* $\lambda \in \mathbb{N}$ *the algorithm has the following specifications:*

i. ***Statistically far:*** $\|\rho_0 - \rho_1\|_1 \geq D(\lambda)$.

ii. ***Computational indistinguishability:*** *For all non-uniform (possibly quantum)* $T(\lambda)$*-time algorithms* $\mathcal{A}$*:*

$$|\Pr(\mathcal{A}(\rho_0) = 1) - \Pr(\mathcal{A}(\rho_1) = 1)| \leq d(\lambda)\,. \tag{8.15}$$

Furthermore, we say that EFI *is a* (D, d)*-EFI for an algorithm* $\mathcal{A}$*, if the computational indistinguishability holds for* $\mathcal{A}$ *without requiring any bound on the runtime of* $\mathcal{A}$*.*

Remark 8.2.1. *When restricted to classical algorithms, EFI pairs with* $D - d \geq 1/\mathsf{poly}(\lambda)$ *and* $T = \mathsf{poly}(\lambda)$ *imply the existence of one-way functions (e.g. see [Gol90, NR06, BDRV19]). The state of the art for the quantum EFI pairs is more restricted. More precisely, an EFI pair with mixed states and* $D^2 - \sqrt{d} \geq O(1)$ *implies quantum bit commitments [BCQ23].*

In this work, we consider the inefficient-verifier one-way state generators.

Definition 8.2.14 (Non-Uniform One-Way State Generators)**.** *Let $T : \mathbb{N} \to \mathbb{R}^+$ and $\theta : \mathbb{N} \to [0,1]$. A (T,θ)-one-way state generator (OWSG) is a tuple of algorithms $\mathsf{G} := (\mathsf{KeyGen}, \mathsf{StateGen}, \mathsf{Ver})$ with the following specification:*

i. ***Key generation:*** $\mathsf{KeyGen}_h(1^\lambda) \to k$*: is a QPT algorithm that given the security parameter 1^n and a $\mathsf{poly}(\lambda)$-size classical advice h, outputs a classical string $k \in \{0,1\}^\lambda$;*

ii. ***State generation:*** $\mathsf{StateGen}(k) \to \rho_k$*, is a QPT algorithm that given a classical string k, outputs a (possibly mixed) m-qubit quantum state $\rho_k \in \mathsf{MS}_m$;*

iii. $\mathsf{Ver}(k,\rho) \in \{\perp, \top\}$*: is a (possibly unbounded) algorithm that given a classical string k and a quantum state ρ outputs either an accept ($\top$) or a reject ($\perp$).*

Further, they satisfy the following properties:

1. ***Correctness:*** *Outputs of the samplers* $(\mathsf{KeyGen}, \mathsf{StateGen})$ *pass the verification with overwhelming probability, i.e.,*

$$\Pr_{\substack{k \leftarrow \mathsf{KeyGen}_h(1^\lambda) \\ \rho_k \leftarrow \mathsf{StateGen}(k)}} (\mathsf{Ver}(k,\rho_k) = 1) \geq 1 - \mathsf{negl}(\lambda)\,. \tag{8.16}$$

2. ***Security:*** *For every non-uniform $T(\lambda)$-time adversary $\mathcal{A}$, and any polynomial t:*

$$\Pr_{\substack{k \leftarrow \mathsf{KeyGen}_h(1^\lambda) \\ \rho_k \leftarrow \mathsf{StateGen}(k) \\ k' \leftarrow \mathcal{A}(\rho_k^{\otimes t}; h)}} \big(\mathsf{Ver}(k',\rho_k) = 1\big) \leq \theta(\lambda)\,. \tag{8.17}$$

Furthermore, we say that G is a θ-OWSG for an algorithm $\mathcal{A}$ if the inequality concerning security (Property 2) holds for $\mathcal{A}$ without requiring any bound on the runtime of $\mathcal{A}$.

We typically refer to OWSGs when $\theta(\lambda) = \mathsf{negl}(\lambda)$, and *weak* OWSG when $T = \mathsf{poly}(\lambda)$ and $\theta = 1 - 1/\lambda^c$ for some constant c. It was shown in [MY24b] that weak OWSGs imply OWSGs.

8.3 Lossy mappings and disguising lemma

Drucker [Dru12] derives a quantitative approach (called the disguising distribution lemma) to measure how much information can be recovered from the output of a compressing mapping about its input, based on the compression size; a distinguishing variant of Fano's inequality. Such mappings are indeed a special type of lossy mappings, an observation upon which Ball, Boyle, Degwekar, Deshpande, Rosen and Vaikuntanathan [BBD+20] build classical cryptographic primitives.

In this section, we focus on variants of lossy mappings and their properties, and extend the disguising lemma from compressing reductions to lossy reductions for a more general class of functions. In our analysis, we consider quantum mappings, but note that all statements hold with respect to randomized functions as well (by considering a classical output).

Classically, a randomized function $R : \{0,1\}^* \to \{0,1\}^*$ is said to be $\ell(\lambda)$-lossy for a class of distributions $X = \{X_\lambda\}_{\lambda\in\mathbb{N}}$ if $I(X_\lambda; R(X_\lambda)) \le \ell(\lambda)$. Below, we also consider general mappings with classical input and quantum output.

Definition 8.3.1 (Lossy Mapping)**.** *Let $\ell : \mathbb{N} \to \mathbb{R}^+$. Let $R : \{0,1\}^* \to \mathsf{MS}_*$ be a quantum mapping. We say that R is $\ell(\lambda)$-lossy for a class of distributions $X = \{X_\lambda\}_{\lambda\in\mathbb{N}}$ over $\{0,1\}^*$, if it holds that*

$$I_q(X_\lambda; R(X_\lambda)) \le \ell(\lambda) \ .$$

For the sake of simplicity, we say that R is $\ell(\lambda)$-lossy, if it is $\ell(\lambda)$-lossy for all distributions.

The results by [Dru12, BBD+20] rely on the lossiness of the mapping for all distributions. Such a condition seems quite strong, in particular, for the multi-variate mappings over m-tuple input. We simplify this condition in two different directions. First, we consider lossy mappings over a particular class of distributions supported on two disjoint sets. These sets will later be YES and NO instances of a promise problem, which allow to show that important classes of reductions are of this form.

Definition 8.3.2 (Splitting Lossy Mapping)**.** *Let $\ell : \mathbb{N} \to \mathbb{R}^+$, $m \in \mathbb{N}$ and $S_0, S_1 \subseteq \{0,1\}^*$ be two disjoint sets. A mapping R is splitting ℓ-lossy supported on (S_0, S_1) if it is ℓ-lossy for the class of distributions $X = (X_1, \ldots, X_m)$ such that for each $i \in [m]$, either $\mathrm{Supp}(X_i) \subseteq S_0$ or $\mathrm{Supp}(X_i) \subseteq S_1$. In other words, $\{X_1, X_2, \cdots, X_m\}$ splits into distributions that are S_0-supported and S_1-supported.*

Remark 8.3.1. *A lossy mapping as per Definition 8.3.1 is also a splitting lossy mapping.*

Later, for the lossy reductions of a problem Π, we choose S_0 and S_1 as the sets Π_{N} and Π_{Y}. Splitting the distribution in such a way allows us to precisely calculate the lossiness of worst-case to average-case reductions.

In the rest of this section, we discuss the generalization of the disguising distribution lemma in [Dru12] and its improvement by [BBD$^+$20]. In both of these results, the lossiness (compression in the former and lossiness in the latter) is considered as in Definition 8.3.1 with respect to all possible input distributions. Instead, we adapt it for splitting lossy maps where the input distribution is uniform over a sparse set. This is obtained by a more refined analysis but yet very similar to those of [Dru12, BBD$^+$20]. Below, we have the main lemma of this section.

Lemma 8.3.3 (Extended Disguising Lemma)**.** *Let λ, m, m_0, m_1 be positive integers such that $m = m_0 + m_1 + 1$, and $R : \{0,1\}^* \to \mathsf{MS}_*$ be any quantum mapping. Further, let $S_0, S_1 \subseteq \{0,1\}^\lambda$ be two disjoint sets, d be a positive integer, $\varepsilon > 0$ be real, and $s(\lambda) := \lceil \lambda \ln 2/(2\varepsilon^2) \rceil$.*

For any choice of positive real $\ell(\lambda)$, if R is splitting $\ell(\lambda)$-lossy for all ds-uniform distributions supported on (S_0, S_1), then there exist two collections of multisets of d elements $K_1, \cdots, K_s$ and $T_1, \cdots, T_s$ respectively contained in S_0 and S_1, such that

- *for any $y \in S_0$, it holds that*

$$\mathop{\mathbb{E}}_{a\sim\mathcal{U}_{[s]},\pi\sim\mathfrak{S}_m}\left[\left\|R(\pi(\mathcal{U}_{K_a}^{\otimes m_0}, y, \mathcal{U}_{T_a}^{\otimes m_1})) - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes(m_0+1)}, \mathcal{U}_{T_a}^{\otimes m_1}\right)\right)\right\|_1\right] \leq \delta_\ell(\lambda) + \frac{2(m+1)}{d+1} + 2\varepsilon \; ; \quad (8.18)$$

- *and for any $y \in S_1$, it holds that*

$$\mathop{\mathbb{E}}_{a\sim\mathcal{U}_{[s]},\pi\sim\mathfrak{S}_m}\left[\left\|R(\pi(\mathcal{U}_{K_a}^{\otimes m_0}, y, \mathcal{U}_{T_a}^{\otimes m_1})) - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m_0}, \mathcal{U}_{T_a}^{\otimes(m_1+1)}\right)\right)\right\|_1\right] \leq \delta_\ell(\lambda) + \frac{2(m+1)}{d+1} + 2\varepsilon \, , \quad (8.19)$$

where

$$\delta_\ell(\lambda) := \min\left\{\sqrt{\frac{\ell(\lambda)\ln 2}{2m}}, 1 - 2^{-\frac{\ell(\lambda)}{m}-2}\right\}. \quad (8.20)$$

Note that the states inside the trace distance are mixed states since the inputs of R are randomized classical distributions.

The proof requires some background definitions and lemmas. Similar to [Dru12, BBD$^+$20], we define distributional stability as follows.

Definition 8.3.4. *Let λ, m, m_0, m_1 be positive integers such that $m = m_0 + m_1 + 1$. For a real $\delta(\lambda) \in [0,1]$, a quantum mapping $R : \{0,1\}^{m\cdot\lambda} \to \mathsf{MS}_*$ is said to be $\delta(\lambda)$-quantumly-distributionally stable ($\delta(\lambda)$-QDS) with respect to two distributions $(\mathcal{D}_0, \mathcal{D}_1)$ over $\{0,1\}^\lambda$ if the following holds:*

$$\mathop{\mathbb{E}}_{y\sim\mathcal{D}_0,\pi\sim\mathfrak{S}_m}\left[\left\|R(\pi(\mathcal{D}_0^{\otimes m_0}, y, \mathcal{D}_1^{\otimes m_1})) - R\Big(\pi\Big(\mathcal{D}_0^{\otimes(m_0+1)}, \mathcal{D}_1^{\otimes m_1}\Big)\Big)\right\|_1\right] \le \delta(\lambda)\,.$$

Note that the order of the pair $(\mathcal{D}_0, \mathcal{D}_1)$ matters. Furthermore, when $m_1 = 0$, we simply say that the mapping is $\delta(\lambda)$-QDS with respect to $\mathcal{D}_0$.

Below, we state an adaptation of [Dru12, Lemma 8.10].

Lemma 8.3.5. *Assume that $R : \{0,1\}^{m\cdot\lambda} \to \mathsf{MS}_*$ satisfies the properties in Lemma 8.3.3 for $m_1 = 0$. Then R is $\delta_\ell(\lambda)$-QDS with respect to any ds-uniform distribution $\mathcal{D}_0$ supported on either S_0 or S_1.*

In the original lemma from [Dru12], compression is used to bound the entropy of the mutual information. However, note that splitting lossiness implies a bound on the mutual information by definition, and any restriction on the input distributions will give a result for the same restricted case.

The following lemma is the generalization of the above one.

Lemma 8.3.6. *Assume that $R : \{0,1\}^{m\cdot\lambda} \to \mathsf{MS}_*$ satisfies the properties in Lemma 8.3.3. Then R is $\delta_\ell(\lambda)$-QDS with respect to any ds-uniform independent distributions $(\mathcal{D}_0, \mathcal{D}_1)$ each supported on either S_0 or S_1.*

Proof. The proof is similar to that of [BBD$^+$20, Proposition B.1]. Let $\pi \in \mathfrak{S}_m$ be a fixed permutation. One can rewrite it as the composition of two partial permutations π_0 and π_1, i.e., $\pi = \pi_0 \circ \pi_1$, such that π_1 only acts on the last m_1 arguments of the input. Let $\rho_\pi(y)$ be as follows

$$\rho_\pi(y) := R\big(\pi\big(\mathcal{D}_0^{\otimes m_0}, y, \mathcal{D}_1^{\otimes m_1}\big)\big)\,. \tag{8.21}$$

For $y, y' \sim \mathcal{D}_0$, two independent random variables, and $\pi \sim \mathfrak{S}_m$, we want to prove that

$$\mathbb{E}_{y,y'\pi}\left[\left\|\rho_\pi(y) - \rho_\pi(y')\right\|_1\right] \le \delta(\lambda)\,. \tag{8.22}$$

Note that it is enough to bound the conditional distributions since

$$\mathbb{E}_{y,y'\pi}\left[\left\|\rho_\pi(y)-\rho_\pi(y')\right\|_1\right]=\mathbb{E}_{\pi_1}\left[\mathbb{E}_{y,y'\pi|\pi_1}\left[\left\|\rho_\pi(y)-\rho_\pi(y')\right\|_1\right]\right], \tag{8.23}$$

by the law of total probability. Let $R'(x_1,x_2,\cdots,x_{m_0+1})$ be the mapping that first samples π then evaluates $R\big(\pi_1\big(x_1,x_2,\cdots,x_{m_0+1},\mathcal{D}_1^{\otimes m_1}\big)\big)$. For any fixed π_1, we show that R' is splitting $\ell(\lambda)$-lossy for all ds-uniform distributions over either S_0 or S_1. Indeed, let $(\mathcal{X}_1,\cdots,\mathcal{X}_{m_0+1})$ be independent ds-uniform random variables with $\mathrm{Supp}(\mathcal{X}_i)\subseteq S_0$ or $\mathrm{Supp}(\mathcal{X}_i)\subseteq S_1$ for each $i\in[m_0+1]$, and $(\mathcal{Z}_1,\cdots,\mathcal{Z}_{m_1})\sim\mathcal{D}_1^{\otimes m_1}$, thus $\mathrm{Supp}(\mathcal{Z}_i)\subseteq\mathrm{Supp}(\mathcal{D}_1)\subseteq S_j$ for all $i\in[m_1]$ and some $j\in\{0,1\}$. By the splitting lossiness of R for any ds-uniform distribution, we can bound the loss of R':

$$\ell(\lambda) \tag{8.24}$$

$$\geq I_q(\pi_1(\mathcal{X}_1,\cdots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\cdots,\mathcal{Z}_{m_1});R(\pi_1(\mathcal{X}_1,\cdots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\cdots,\mathcal{Z}_{m_1}))) \tag{8.25}$$

$$=I_q(\mathcal{X}_1,\cdots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\cdots,\mathcal{Z}_{m_1};R(\pi_1(\mathcal{X}_1,\cdots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\cdots,\mathcal{Z}_{m_1}))) \tag{8.26}$$

$$\geq I_q(\mathcal{X}_1,\cdots,\mathcal{X}_{m_0+1};R(\pi_1(\mathcal{X}_1,\cdots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\cdots,\mathcal{Z}_{m_1}))). \tag{8.27}$$

Finally, by Lemma 8.3.5 a splitting lossy map must also be $\delta_\ell(\lambda)$-QSD, thus

$$\mathbb{E}_{y,y'\pi|\pi_1}\left[\left\|\rho_\pi(y)-\rho_\pi(y')\right\|_1\right] \tag{8.28}$$

$$=\mathbb{E}_{y,y'\pi|\pi_1}\left[\left\|R\big(\pi\big(\mathcal{D}_0^{\otimes m_0},y,\mathcal{D}_1^{\otimes m_1}\big)\big)-R\big(\pi\big(\mathcal{D}_0^{\otimes m_0},y',\mathcal{D}_1^{\otimes m_1}\big)\big)\right\|_1\right] \tag{8.29}$$

$$=\mathbb{E}_{y,y'\pi_0}\left[\left\|R'\big(\mathcal{D}_0^{\otimes m_0},y\big)-R'\big(\mathcal{D}_0^{\otimes m_0},y'\big)\right\|_1\right] \tag{8.30}$$

$$\leq\delta_\ell(\lambda)\,. \tag{8.31}$$

□

If a mapping is distributionally stable with respect to a pair of distributions, then one can "sparsify" the distributions while nearly keeping the stability.

Lemma 8.3.7. *Let $\lambda,m,m_0,m_1,\ell(\lambda),S_0,S_1,R,d$ and $\delta_\ell(\lambda)$ be as defined in Lemma 8.3.3. Let $\mathcal{D}_0$ and $\mathcal{D}_1$ be two independent distributions with supports over S_0 and S_1, respectively. Let $\{x_i^{(0)}\}_{i\in[d+1]}$ and $\{x_i^{(1)}\}_{i\in[d+1]}$ be independent samples from $\mathcal{D}_0$ and $\mathcal{D}_1$, respectively. For each $j\in\{0,1\}$,* *let $y_j^* := x_{i^*}^{(j)}$ be uniformly chosen from $\{x_i^{(j)}\}_{i\in[d+1]}$ and let $\widehat{\mathcal{D}}_j$ be the uniform distribution over the multiset $\{x_i^{(j)}\}_{i\in[d+1]\setminus\{i^*\}}$. Then it holds that*

$$\mathop{\mathbb{E}}_{\pi\sim\mathfrak{S}_m}\left[\left\|R\Big(\pi\Big(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\Big)\Big) - R\Big(\pi\Big(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\Big)\Big)\right\|_1\right] \leq \delta_\ell(\lambda) + \frac{2m_0+1}{d+1}\,, \tag{8.32}$$

$$\mathop{\mathbb{E}}_{\pi\sim\mathfrak{S}_m}\left[\left\|R\Big(\pi\Big(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_1^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\Big)\Big) - R\Big(\pi\Big(\widehat{\mathcal{D}}_0^{\otimes m_0}, \widehat{\mathcal{D}}_1^{\otimes (m_1+1)}\Big)\Big)\right\|_1\right] \leq \delta_\ell(\lambda) + \frac{2m_1+1}{d+1}\,. \tag{8.33}$$

Proof. We prove the first statement. The other one is implied similarly. Let $\widetilde{\mathcal{D}}_0$ denote the uniform distribution over $\{x_i^{(0)}\}_{i\in[d+1]}$. For any fixed set of of multisets as above and any choice of permutation π and quantum mapping R, we have

$$\begin{aligned}
&\left\|R\Big(\pi\Big(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\Big)\Big) - R\Big(\pi\Big(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\Big)\Big)\right\|_1 && (8.34)\\
&\leq \left\|\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)} \otimes \widehat{\mathcal{D}}_1^{\otimes m_1} - \widehat{\mathcal{D}}_0^{\otimes (m_0+1)} \otimes \widehat{\mathcal{D}}_1^{\otimes m_1}\right\|_1 && (8.35)\\
&\leq \left\|\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)} - \widehat{\mathcal{D}}_0^{\otimes (m_0+1)}\right\|_1 && (8.36)\\
&\leq (m_0+1)\left\|\widetilde{\mathcal{D}}_0 - \widehat{\mathcal{D}}_0\right\|_1\,, && (8.37)
\end{aligned}$$

where we used the quantum data processing inequality for the first two upper bounds, and the tensor product property for the last one. Since both $\widehat{\mathcal{D}}_0$ and $\widetilde{\mathcal{D}}_0$ are classical, their trace distance coincides with their statistical distance. Therefore, we have

$$\begin{aligned}
\left\|\widetilde{\mathcal{D}}_0 - \widehat{\mathcal{D}}_0\right\|_1 &= \frac{1}{2}\sum_{x\in\{x_i^{(0)}\}_{i\in[d+1]}} |\Pr_{\widetilde{\mathcal{D}}_0}(x) - \Pr_{\widehat{\mathcal{D}}_0}(x)| && (8.38)\\
&= \frac{1}{2(d+1)} + \frac{1}{2}\sum_{x\in\{x_i^{(0)}\}_{i\in[d+1]\setminus\{i^*\}}} \left|\frac{1}{d+1} - \frac{1}{d}\right| && (8.39)\\
&= \frac{1}{d+1}\,. && (8.40)
\end{aligned}$$

Similarly, it holds that

$$\left\| R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) \right\|_1 \leq \frac{m_0}{d+1} \; . \tag{8.41}$$

From the triangle inequality, it follows that

$$\begin{aligned}
&\left\| R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) \right\|_1 && (8.42)\\
&\leq \left\| R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) \right\|_1 \\
&\quad + \left\| R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) \right\|_1 && (8.43)\\
&\quad\quad + \left\| R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) \right\|_1 \\
&< \left\| R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) \right\|_1 + \frac{2m_0+1}{d+1} \; . && (8.44)
\end{aligned}$$

Recall that R is splitting $\ell(\lambda)$-lossy with respect to all ds-uniform distributions supported on (S_0, S_1). Therefore, by Lemma 8.3.6 it is $\delta_\ell(\lambda)$-QSD with respect to all ds-uniform pair of distributions each supported on either S_0 or S_1, including $(\widetilde{\mathcal{D}}_0, \widehat{\mathcal{D}}_1)$. Finally, by taking the expectation from both sides above with respect to π, and using the fact that R is $\delta_\ell(\lambda)$-QSD with respect to $(\widetilde{\mathcal{D}}_0, \widehat{\mathcal{D}}_1)$, one obtains the claimed upper bound. □

Proof of Lemma 8.3.3. Consider the following two-player, simultaneous, zero-sum game:

- Player 1: chooses a pair of multisets $K \subseteq S_0$ and $T \subseteq S_1$, each of size d.
- Player 2: chooses an element $y \in S_0 \cup S_1$
- Payoff: if $y \in S_0$, Player 2 gains

  $$\mathop{\mathbb{E}}_{\pi \sim \mathfrak{S}_m}\left[\left\| R(\pi(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1})) - R\left(\pi\left(\mathcal{U}_K^{\otimes (m_0+1)}, \mathcal{U}_T^{\otimes m_1}\right)\right) \right\|_1\right] \; , \tag{8.45}$$

  otherwise, Player 2 gains

  $$\mathop{\mathbb{E}}_{\pi \sim \mathfrak{S}_m}\left[\left\| R(\pi(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1})) - R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, \mathcal{U}_T^{\otimes (m_1+1)}\right)\right) \right\|_1\right] \; . \tag{8.46}$$

Consider a ds-uniform strategy for Player 2, i.e. a distribution $\mathcal{Y}$ of y that is uniform over a multiset of pure strategies of size ds. We explain a strategy $(\mathcal{K}, \mathcal{T})$ for Player 1 that bounds the expected payoff. Player 1 chooses K by sampling d independent instances of the restriction of $\mathcal{Y}$ to S_0, and chooses T by sampling d independent instances of the restriction of $\mathcal{Y}$ to S_1. The expected payoff E is

$$\begin{aligned}\Pr_{y\sim\mathcal{Y}}(y\in S_0)\mathbb{E}_{\pi,K,T}\Big[\Big\|R(\pi(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}))\\ -R\Big(\pi\Big(\mathcal{U}_K^{\otimes(m_0+1)}, \mathcal{U}_T^{\otimes m_1}\Big)\Big)\Big\|_1\Big|y\in S_0\Big]\\ +\Pr_{y\sim\mathcal{Y}}(y\in S_1)\mathbb{E}_{\pi,K,T}\Big[\Big\|R(\pi(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}))\\ -R\Big(\pi\Big(\mathcal{U}_K^{\otimes m_0}, \mathcal{U}_T^{\otimes(m_1+1)}\Big)\Big)\Big\|_1\Big|y\in S_1\Big]\,. \qquad (8.47)\end{aligned}$$

Let $x_1^{(0)}, x_2^{(0)}, \cdots, x_{d+1}^{(0)}$ and $x_1^{(1)}, x_2^{(1)}, \cdots, x_{d+1}^{(1)}$ be $d+1$ independent samples from $\mathcal{Y}|_{S_0}$ and $\mathcal{Y}|_{S_1}$, respectively. Sample $i^* \xleftarrow{\$} [d+1]$ and for $j \in \{0,1\}$, let $y_j^* := x_{i^*}^{(j)}$. Let $\widehat{\mathcal{Y}}_0$ and $\widehat{\mathcal{Y}}_1$ be the uniform distributions over the multisets $\{x_i^{(0)}\}_{i\in[d+1]\setminus\{i^*\}}$ and $\{x_i^{(1)}\}_{i\in[d+1]\setminus\{i^*\}}$, respectively. For $j \in \{0,1\}$, we have that $(y_j^*, \widehat{\mathcal{Y}}_0, \widehat{\mathcal{Y}}_1) \sim (\mathcal{Y}|_{S_j}, \mathcal{K}, \mathcal{T})$. Then, by Lemma 8.3.7, we have

$$\begin{aligned}\mathbb{E}_\pi\Big[\Big\|R\Big(\pi\Big(\widehat{\mathcal{Y}}_0^{\otimes m_0}, y, \widehat{\mathcal{Y}}_1^{\otimes m_1}\Big)\Big) - R\Big(\pi\Big(\widehat{\mathcal{Y}}_0^{\otimes(m_0+1)}, \widehat{\mathcal{Y}}_1^{\otimes m_1}\Big)\Big)\Big\|_1 \,\Big|\, y\in S_0\Big]\\ \le \delta_\ell(\lambda) + \frac{2m_0+1}{d+1}\,, \qquad (8.48)\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}_\pi\Big[\Big\|R\Big(\pi\Big(\widehat{\mathcal{Y}}_0^{\otimes m_0}, y, \widehat{\mathcal{Y}}_1^{\otimes m_1}\Big)\Big) - R\Big(\pi\Big(\widehat{\mathcal{Y}}_0^{\otimes m_0}, \widehat{\mathcal{Y}}_1^{\otimes(m_1+1)}\Big)\Big)\Big\|_1 \,\Big|\, y\in S_1\Big]\\ \le \delta_\ell(\lambda) + \frac{2m_1+1}{d+1}\,. \qquad (8.49)\end{aligned}$$

Therefore, we obtain $E \le \delta_\ell(\lambda) + 2(m+1)/(d+1)$.

Above, we showed that for every ds-uniform strategy for Player 2, there exists a strategy for Player 1 that bounds the expected payoff by $\delta_\ell(\lambda)+2(m+1)/(d+1)$. Let $\mathbf{M} := [\mathbf{M}_{ij}]_{i,j}$ be the matrix such that $\mathbf{M}_{ij}$ corresponds to the payoff when Player 1 outputs i and Player 2 outputs j. By Lemma 8.2.2, we have

$$\begin{aligned}\delta_\ell(\lambda) + 2(m+1)/(d+1) &\ge \max_{\mathcal{Q}\in\mathfrak{Q}_{ds}} \min_i \mathbb{E}_{j\sim\mathcal{Q}}[\mathbf{M}_{ij}] && (8.50)\\ &\ge \omega(\mathbf{M}) - \varepsilon(\mathbf{M}_{\max} - \mathbf{M}_{\min}) && (8.51)\\ &\ge \omega(\mathbf{M}) - \varepsilon\,, && (8.52)\end{aligned}$$

where $\mathfrak{Q}_{ds}$ is the set of all ds-uniform strategies for Player 2. It follows that $\omega(\mathbf{M}) \leq \delta_\ell(\lambda) + 2(m+1)/(d+1) + \varepsilon$.

Now we use Lemma 8.2.2 the other way around. In fact, the number of possible choices for Player 1 is $|S_0 \cup S_1| \leq 2^\lambda$. Therefore, Lemma 8.2.2 asserts that there exists an s-uniform strategy for Player 2 such that for any possible mixed strategy for Player 1, the expected payoff is at most ε-far from the value of the game $\omega(\mathbf{M})$. In other words, for this particular strategy of Player 1, the expected payoff is always at most

$$\omega(\mathbf{M}) + \varepsilon \leq \delta_\ell(\lambda) + 2(m+1)/(d+1) + 2\varepsilon \,. \tag{8.53}$$

Recall that a $s(\lambda)$-uniform strategy is, by definition, a uniformly sampled element from a size-$s(\lambda)$ multiset of choices of the player. Note that Player 1 chooses a pair (K, T). Therefore, this strategy is essentially a uniform distribution over some multiset $\{(K_1, T_1), \cdots, (K_s, T_s)\}$, which concludes the proof.

□

8.4 Mildly-lossy problems

In this section, we first put forward a new abstraction, which we call an f-distinguisher reduction that implies definitions of f-reductions as well as Karp and non-adaptive Turing reductions. Then, by considering the lossiness property (as defined in Section 8.3), we introduce mildly-lossy problems which will be the core of our analysis in the subsequent sections. In particular, in Section 8.6 we construct quantum cryptographic primitives from mildly-lossy problems we will show in Section 8.7 that if a problem admits a worst-case to average-case f-distinguisher reductions then it is mildly lossy (under certain conditions).

The analysis of the following two Sections 8.4 and 8.5 applies to both classical and quantum reductions. For the sake of simplicity and generality, we only refer to quantum reductions and we explicitly highlight the distinction when necessary.

8.4.1 f-Distinguisher Reductions

A Karp decision-to-decision reduction R from Π to Σ has the following property: $\chi_\Pi(x) = 1$ if and only if $\chi_\Sigma(R(x))$ (up to some error). In our work, the target problem Σ is not restricted and does not play any role. Therefore, we consider the following more general notion: a mapping R is a reduction if there exists a (possibly unbounded) distinguisher $\mathcal{D}$ that can tell $R(x)$ and $R(x')$ apart, when $\chi_\Pi(x) \neq \chi_\Pi(x')$ (up to some error). A reduction is therefore a mapping that preserves the distinguishing power of the unbounded algorithm[1]. In other words, it preserves some information about the inputs. When the reduction is to a search problem, there must also exist an inverting algorithm that given x and the solution (or witness) of $R(x)$, outputs $\chi_\Pi(x)$. To include such reductions, we generalize this definition once more by allowing the distinguisher to have one and only one of the instances x or x'. To see how this helps, we give an example: the reduction from ParamSat to MaxSat. In ParamSat, an instance $x := (\varphi, k)$, with φ a CNF formula and k an integer, is a YES instance if and only if at least k clauses of φ are satisfiable. The MaxSat problem asks to find an assignment that satisfies the maximum number of clauses. Consider the decision-to-search reduction as follows: given an instance $x := (\varphi, k)$ of ParamSat, the output of the reduction is φ. By having k and an assignment w_φ satisfying the maximum number of clauses of φ (the solution of φ as a MaxSat instance), the reduction computes $\chi_{\text{ParamSat}}(x)$ by comparing k and the number of satisfied clauses by w_φ. Note that it is necessary for the inverting algorithm to know k.

In this subsection, we introduce the generalized distinguisher reductions.

Definition 8.4.1 (f-Distinguisher Reduction)**.** *Let λ, m be positive integers, and $\mu : \mathbb{N} \to [0,1]$ be a function of λ. Let $f : \{0,1\}^m \to \{0,1\}$, and Π be a promise problem. A (μ, f^m)-distinguisher reduction for Π is a quantum mapping $R : \{0,1\}^* \to \mathsf{MS}_*$, for which there exists a distinguisher $\mathcal{D}$, such that for all $(x_1, \cdots, x_m)$ and $(x'_1, \cdots, x'_m)$ in $((\Pi_Y \cup \Pi_N) \cap \{0,1\}^\lambda)^m$ where $f(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m)) \neq f(\chi_\Pi(x'_1), \cdots, \chi_\Pi(x'_m))$, we have*

$$\mathbb{E}_{i \sim \mathcal{U}_{[m]}} \big| \Pr(1 \leftarrow \mathcal{D}(h_i, R(x_1, \cdots, x_m))) - \Pr\big(1 \leftarrow \mathcal{D}(h_i, R(x'_1, \cdots, x'_m))\big) \big| \\ \geq 1 - 2\mu(\lambda) ,$$

where $h_i := (x_i, \{\chi_\Pi(x_j)\}_j, \{\chi_\Pi(x'_j)\}_j)$. We call μ the error of the reduction.

Drucker [Dru12, Definition 8.2] defines an f-compression reduction for a promise problem Π in a somewhat similar fashion we define f-distinguisher reductions: as a mapping that sends instances $x_1, \cdots, x_m$ of size λ to a

[1] Note that an unbounded algorithm can always distinguish YES and NO instances of a problem by simply solving them.

quantum state ρ, such that there exists a binary measurement $\mathcal{M}$ (not necessarily efficient) that outputs $f(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m))$ with probability more than $1 - \mu(\lambda)$. Recall his definition.

Definition 8.4.2 (f-Reduction)**.** *Let λ, m be positive integers, and $\mu : \mathbb{N} \to [0,1]$ be a function of λ. Let $f : \{0,1\}^m \to \{0,1\}$, and Π be a promise problem. A (μ, f^m)-reduction for Π is a mapping $R : \{0,1\}^{m\cdot\lambda} \to S$, where $S = \{0,1\}^*$ (classical) or $S = \mathsf{MS}_*$ (quantum), for which there exists a family of unbounded algorithms $\{\mathcal{M}_k\}_{k\in\mathbb{N}}$, such that for all $(x_1, \cdots, x_m) \in ((\Pi_Y \cup \Pi_N) \cap \{0,1\}^\lambda)^m$,*

$$\Pr(\mathcal{M}(R(x_1, \cdots, x_m)) = f(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m))) \geq 1 - \mu(\lambda)\,,$$

where the probability is taken over the randomness of R and $\mathcal{M}$. We call μ the error of the reduction.

In the following, we will show that f-reductions are special cases of f-distinguisher reductions (per Definition 8.4.1) when the hint h_i is set to be empty.

Lemma 8.4.3. *Let $f : \{0,1\}^m \to \{0,1\}$, and Π be a promise problem. If R is a (μ, f^m)-reduction for Π, then R is also a (μ, f^m)-distinguisher reduction for Π.*

Proof. Recall that for an f-reduction there exists a binary quantum measurement $\mathcal{M}$ such that for every $(x_1, \ldots, x_m) \in ((\Pi_Y \cup \Pi_N) \cap \{0,1\}^\lambda)^m$ decides the f-reduction:

$$\Pr(\mathcal{M}(R(x_1, \cdots, x_m)) = f(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m))) \geq 1 - \mu(\lambda)\,, \tag{8.54}$$

which implies that $\mathcal{M}$ can distinguish $R(x_1, \cdots, x_m)$ from $R(x'_1, \cdots, x'_m)$ for any two inputs $(x_1, \ldots, x_m), (x'_1, \ldots, x'_m) \in ((\Pi_Y \cup \Pi_N) \cap \{0,1\}^\lambda)^m$ with probability at least $1 - 2\mu(\lambda)$. Therefore, there exists an unbounded distinguisher $\mathcal{D}$ such that for h_i per Definition 8.4.1, we have

$$\begin{aligned}
&\mathbb{E}_{i\sim\mathcal{U}_{[m]}}\big|\Pr[1 \leftarrow \mathcal{D}(h_i, R(x_1, \cdots, x_m))] - \Pr\big[1 \leftarrow \mathcal{D}(h_i, R(x'_1, \cdots, x'_m))\big]\big| && (8.55)\\
&\geq \big|\Pr(1 \leftarrow \mathcal{M}(R(x_1, \cdots, x_m))) - \Pr\big(1 \leftarrow \mathcal{M}(R(x'_1, \cdots, x'_m))\big)\big| && (8.56)\\
&\geq 1 - 2\mu(\lambda)\,, && (8.57)
\end{aligned}$$

where for the first inequality we used the fact that revealing more information to the distinguisher does not decrease its advantage. □

8.4.2 Mildly-Lossy Problems

To analyse the lossiness of f-distinguisher reductions, we fix the set of functions f to those ones that are invariant under permuting their inputs.

Definition 8.4.4 (Permutation-Invariant)**.** *A Boolean function $f : \{0,1\}^m \to \{0,1\}$ is called permutation-invariant if for every permutation $\pi \in \mathfrak{S}_m$, it holds that $f(\pi(b_1, b_2, \cdots, b_m)) = f(b_1, b_2, \cdots, b_m)$.*

Note that these are sometimes referred to in the literature as symmetric Boolean functions.

This set of functions is of great interest. The functions And, Or, and Maj that were considered in [Dru12, BBD+20] are all non-constant permutation-invarant. Moreover, the (non-monotone) functions Parity and Mod$_k$ are of this type as well as Threshold$_k$.

We use the following technical lemma about non-constant permutation-invariant functions.

Lemma 8.4.5. *Let $f : \{0,1\}^m \to \{0,1\}$ be a non-constant permutation-invariant function. Then there exists an integer $1 \leq p \leq m$ such that*

$$f(\underbrace{1,1,\cdots,1}_{p-1},0,0,\cdots,0) = 0\,, \quad \textit{and} \quad f(\underbrace{1,1,\cdots,1}_{p},0,0,\cdots,0) = 1\ .$$

We let $p(f)$ denote the minimum choice of such an integer.

Proof. The set $\{0,1\}^m$ can be partitioned into $m+1$ equivalence classes where each class consists of strings with the same number of 1's. We note that the result of a permutation on an input falls in the same equivalence class. Therefore, since the function is permutation-invariant, then the evaluation of f over each input is determined by its class. Because the function is non-constant, there must exist two consecutive classes (the classes can be ordered by the number of 1's that they represent) with different evaluation under f. This completes the proof. □

Finally, we introduce the notion of *mildly-lossy problems* which are promise problems that admit lossy f-distinguisher reductions where f is a non-constant permutation-invariant function.

Definition 8.4.6 (Mildly-Lossy Problems)**.** *Let λ, m be positive integers, γ be a positive real and T, μ, α be functions on λ such that $\mu(\lambda) \in [0, 1/2)$. A promise problem Π is said to be $(T, \mu, f^m, \alpha, \gamma)$-mildly-lossy if there exists a non-uniform (μ, f^m)-distinguisher reduction R (per Definition 8.4.1) for Π with the following properties:*

i. f is some non-constant permutation-invariant function $f : \{0,1\}^m \to \{0,1\}$,

ii. the reduction R runs in time $T(\lambda)$,

iii. R is splitting $m\alpha(\lambda)$-lossy (per Definition 8.3.2) supported on (Π_Y, Π_N), for all pairwise independent $(2^9 m\lambda/\gamma^3)$-uniform distributions over λ-bit strings.

We interchangeably say that the reduction R as above is mildly-lossy.

For simplicity of notation, we will redefine δ_ℓ from Lemma 8.3.3 for $m\alpha(\lambda)$-lossy maps.

Definition 8.4.7. *We let $\delta\colon : \mathbb{N}^+ \to \mathbb{R}$ be the following function*

$$\delta(\lambda) := \min\left\{\sqrt{\frac{\lambda \ln 2}{2}}, 1 - 2^{-\lambda-2}\right\}.$$

8.5 Zero-knowledge from mildly-lossy problems

In this section, we show that mildly-lossy problems admit Karp reductions to the quantum state distinguishability problem. We provide a fine-grained analysis. When restricted to polynomial-time AND-compression reductions, this recreates the result of Drucker [Dru12, Theorem 8.14]: roughly, if a promise problem Π has a (quantum) polynomial-time AND-compression reduction, then Π must belong to SZK (resp., QSZK). Similar statements holds for the AND- or MAJ-lossy reductions (see [BBD+20]). We note that our result holds for any non-constant permutation-invariant function, requires a less restricted notion of lossiness, and allows superpolynomial-time reductions.

Theorem 8.5.1. *Let Π be $(T, \mu, f^m, \alpha, \gamma)$-mildly-lossy. Assume that $\theta_{\mathsf{szk}} := (1-2\mu(\lambda))^2/(\delta(\alpha(\lambda))+\gamma) > 1$, with δ as in Definition 8.4.7. Then Π reduces to a problem in QSZK in time $O((T+m^2\lambda)/(\gamma \log \theta_{\mathsf{szk}}))$ and with a classical advice of size $4m\lambda/\gamma$ as described in Protocol 24. Moreover, the reduction is deterministic (but non-uniform) and Π reduces to SZK if Π is lossy with respect to a classical reduction.*

Remark 8.5.1 (Input-output type of the circuits)**.** *Consider the two circuits $(\widehat{C}_0, \widehat{C}_1)$ in Protocol 24, Lines 3 and 4. When R is a randomized reduction, the two circuits are also randomized. Part of their randomness input*

Protocol 24 Reduction $\mathcal{F}$ from Π to $\mathrm{QSD}_{1/4,3/4}$.

Parameters: $\lambda, m, \mu, f, \alpha, \gamma, R, \Pi$ as in Definition 8.4.6. Further $S_0 := \Pi_N \cap \{0,1\}^\lambda$, $S_1 := \Pi_Y \cap \{0,1\}^\lambda$,

$$\varepsilon := \frac{\gamma}{4}, \quad d := \left\lceil \frac{m+1}{\varepsilon} \right\rceil, \; s := \left\lceil \frac{\lambda \ln 2}{2\varepsilon^2} \right\rceil, \tag{8.58}$$

and $K_1, \cdots K_s, T_1, \cdots, T_s$ as in Lemma 8.3.3.

Input: An instance $y \in \{0,1\}^\lambda$.

Advice: $p := p(f)$ as in Lemma 8.4.5, $b_Y, b_N \in \{0,1\}$ respectively representing whether $\Pi_Y \cap \{0,1\}^\lambda$ and $\Pi_N \cap \{0,1\}^\lambda$ are empty. K_a, T_a, π for some uniformly chosen $a \in [s]$ and $\pi \in \mathfrak{S}_m$.

Output: A pair of circuits (C_0, C_1).

```
If $b_N = 1$, return any two circuits $(Y_0, Y_1)$ where $\|Y_0|\mathbf{0}\rangle - Y_1|\mathbf{0}\rangle\|_1 \leq 1/4$.
If $b_Y = 1$, return any two circuits $(N_0, N_1)$ where $\|N_0|\mathbf{0}\rangle - N_1|\mathbf{0}\rangle\|_1 \geq 3/4$.
Let $\widehat{C}_0$ be the following circuit: it samples $\widetilde{x} \sim \left(\mathcal{U}_{K_a}^{\otimes m-p+1}, \mathcal{U}_{T_a}^{\otimes p-1}\right)$, then it outputs $R(\pi(\widetilde{x}))$.
Let $\widehat{C}_1$ be the following circuit: it samples $\widetilde{x} \sim \left(\mathcal{U}_{K_a}^{\otimes m-p}, y, \mathcal{U}_{T_a}^{\otimes p-1}\right)$, then it outputs $R(\pi(\widetilde{x}))$.
Compute $(C_0, C_1) \leftarrow \mathsf{Polarize}(\widehat{C}_0, \widehat{C}_1, 1^2)$ (recall Lemma 8.2.11).
Return $(C_0, C_1)$.
```

is used to sample $\widetilde{x}$ and the other part is fed to R. Let κ be the size of the total randomness. For $r \in \{0,1\}^\kappa$ and any $b \in \{0,1\}$, we let $\widehat{C}_b(r)$ denote the outcome of $\widehat{C}_b$ given the randomness r. On the other hand, when R is quantum, the circuits will be mixed algorithms; classical randomness is required for sampling $\widetilde{x}$. Let κ' be the size of total randomness.[2] *For any $r \in \{0,1\}^{\kappa'}$ and any $b \in \{0,1\}$, we let the mixed outcome of $\widehat{C}_b$ be $\widehat{C}_b|r, \mathbf{0}\rangle$ where $|\mathbf{0}\rangle$ is some appropriate-size ancilla, emphasizing its mixed classical-quantum nature. When it is not relevant, we drop the dependency on r for simplification.*

Proof of Theorem 8.5.1. In the following, we assume that R is quantum. The classical case is similar with the only difference being the type of the inputs and outputs of $(\widehat{C}_0, \widehat{C}_1)$.

[2] Note that κ and κ' are possibly different depending on how much classical randomness R requires.

Consider the case $y \in \Pi_Y$. We bound the ℓ_1 distance (per Definition 8.2.8) of the outcomes of $\widehat{C}_0$ and $\widehat{C}_1$ from below. Sample a uniform coin $b \leftarrow \{0,1\}$, and let $z = \widehat{C}_b|r, \mathbf{0}\rangle$ where r follows the uniform distribution. We drop the dependency on r for simplification. Let $\mathcal{A}$ be a (possibly unbounded) distinguisher that takes z as input and guesses which circuit ($\widehat{C}_0$ or $\widehat{C}_1$) is used to compute z. Let $\mathcal{A}$ be the quantum distinguisher of the (μ, f^m)-distinguisher reduction (that comes from Definition 8.4.6) for Π. On the one hand, if z is computed by $\widehat{C}_0$, we have that $\widetilde{x} := (x_1, \cdots, x_m) \sim \left(\mathcal{U}_{K_a}^{\otimes m-p+1}, \mathcal{U}_{T_a}^{\otimes p-1}\right)$ with $K_a \subseteq \Pi_N \cap \{0,1\}^\lambda$ and $T_a \subseteq \Pi_Y \cap \{0,1\}^\lambda$. Then, since $\widetilde{x}$ contains $p-1$ YES instances by Lemma 8.4.5, for any $\pi \in \mathfrak{S}_m$, we have

$$f(\pi(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m))) = 0\,. \tag{8.59}$$

On the other hand, if z is computed by $\widehat{C}_1$, we have that $\widetilde{x}$ contains one more YES instance $y \in \Pi_Y \cap \{0,1\}^\lambda$, therefore,

$$f(\pi(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m))) = 1\,. \tag{8.60}$$

Moreover, revealing π with the description of the circuits does not decrease the success probability of the distinguisher, thus we can lower bound the trace distance between the circuits for YEs instances by the quantum f-distinguishability of the reduction

$$\begin{aligned}
&\|\widehat{C}_0|\mathbf{0}\rangle - \widehat{C}_1|\mathbf{0}\rangle\|_1 && (8.61)\\
&\geq \mathbb{E}_{i \sim \mathcal{U}_{[m]}} \left|\Pr\left[1 \leftarrow \mathcal{D}(x_i, \widehat{C}_0|\mathbf{0}\rangle)\right] - \Pr\left[1 \leftarrow \mathcal{D}(x_i, \widehat{C}_1|\mathbf{0}\rangle)\right]\right| && (8.62)\\
&\geq 1 - 2\mu(\lambda)\,. && (8.63)
\end{aligned}$$

Now, we discuss the case of $y \in \Pi_N$. We consider a modification of the distinguishing game where the random variables a and π are also given to the distinguisher. Revealing a, π along with z does not decrease the success probability of the distinguisher, thus we can bound the original distinguishing probability by the distinguishing probability of the new task. It holds that

$$\begin{aligned}
&\|\widehat{C}_0|\mathbf{0}\rangle - \widehat{C}_1|\mathbf{0}\rangle\|_1\\
&\quad \leq \left\| R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p+1}, \mathcal{U}_{T_a}^{\otimes p-1}\right)\right) - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p}, y, \mathcal{U}_{T_a}^{\otimes p-1}\right)\right)\right\|_1, \quad (8.64)
\end{aligned}$$

By taking the expectation over a and π, we have

$$\begin{aligned}\|\widehat{C}_0|\mathbf{0}\rangle - \widehat{C}_1|\mathbf{0}\rangle\|_1 \\ &\leq \mathbb{E}_{a\sim\mathcal{U}_{[s]},\pi\sim\mathfrak{S}_m}\Bigg[\Bigg\|R\Big(\pi\Big(\mathcal{U}_{K_a}^{\otimes m-p+1},\mathcal{U}_{T_a}^{\otimes p-1}\Big)\Big) \\ &\qquad - R\Big(\pi\Big(\mathcal{U}_{K_a}^{\otimes m-p}, y, \mathcal{U}_{T_a}^{\otimes p-1}\Big)\Big)\Bigg\|_1\Bigg]. \qquad (8.65)\end{aligned}$$

By our choice of ε, d, s, $K_1, \ldots, K_s, T_1, \cdots, T_s$ and Lemma 8.3.3, we conclude that

$$\|\widehat{C}_0|\mathbf{0}\rangle - \widehat{C}_1|\mathbf{0}\rangle\|_1 \leq \delta(\alpha(\lambda)) + \frac{2(m-p+1)}{d+1} + 2\varepsilon \leq \delta(\alpha(\lambda)) + \gamma\,. \qquad (8.66)$$

$$(8.67)$$

Above, we proved that $(\widehat{C}_0, \widehat{C}_1)$ is an instance of $\mathrm{QSD}_{\delta+\gamma,1-2\mu}$ of size $(T+m^2\lambda)/\gamma$. By assumption, we have $\theta_{\mathsf{szk}} = (1-2\mu(\lambda))^2/(\delta(\alpha(\lambda))+\gamma)$. Therefore, the runtime of $\mathsf{Polarize}(\widehat{C}_0, \widehat{C}_1, 1^2)$ and its output size are both of $O((T+m^2\lambda)/(\gamma\log\theta_{\mathsf{szk}}))$ according to Lemma 8.2.11. □

8.6 Cryptography from mildly-lossy problems

In this section we discuss how mildly-lossy problems can be used to build cryptographic primitives. In Theorem 8.6.1, we construct EFI schemes and in Theorem 8.6.2 OWSGs. The statements allow both classical reductions and quantum reductions.

Our analysis varies based on how lossy $\alpha(\lambda)$ of a reduction a problem admits. When $\alpha(\lambda)$ is small (per Remark 8.6.1), one can leverage Theorem 8.5.1 to build (quantum) EFI pairs. We immediately obtain one-way functions from EFIs (or quantum bit commitments if the reduction is quantum), by taking into account the known constructions from EFI schemes (see Remark 8.2.1). In Theorem 8.6.2, we explain how to obtain one-way state generators when $\alpha(\lambda)$ is large. Finally, we note that the latter does imply one-way functions, too.

8.6.1 EFIs from mildly-lossy problems

Theorem 8.6.1. *Let Π be $(T,\mu,f^m,\alpha,\gamma)$-mildly-lossy. Assume that $\theta_{\mathsf{efi}} := (1-2\mu(\lambda)) - 3(\delta(\alpha(\lambda))+\gamma) > 0$, with δ as in Definition 8.4.7. Then there exists an algorithm* EFI *that runs in $O(T+m^2\lambda\gamma^{-1})$ and an oracle algorithm $\mathcal{C}$,* *such that for any algorithm $\mathcal{A}$ one and only one of the following statements holds:*

i. $\mathcal{C}^{\mathcal{A}}$ *solves* $\Pi \cap \{0,1\}^{\lambda}$ *in time* $O((T + m^2\lambda\gamma^{-1})\theta_{\mathsf{efi}}^{-2})$ *with* $O(\theta_{\mathsf{efi}}^{-2})$ *queries to* $\mathcal{A}$*,*

ii. EFI *is* $(1 - 2\mu(\lambda), 1 - 2\mu(\lambda) - \theta_{\mathsf{efi}}/2)$*-EFI for* $\mathcal{A}$*.*

Moreover, if the mildly-lossy reduction of Π *is classical, then the algorithm* EFI *is also classical.*

Remark 8.6.1. *From the conditions of Theorem 8.6.1, it must hold that* $\delta < 1/3$*, therefore,* $\alpha(\lambda)$ *must be small. Most notably, the statement does not include perfect* 1*-mildly-lossy reductions. However, this can be overcome as follows: Let* R *be* 1*-mildly-lossy and perfect. Consider the new reduction* R' *that with probability* 0.35 *randomly outputs a YES or a NO instance of the target language (note that instance can be given as advice). Otherwise, it applies* R*. The new reduction is* 0.35*-mildly-lossy with error* 0.375 *which satisfies the condition* $(1 - 2\mu(\lambda)) - 3(\delta(\alpha(\lambda)) + \gamma) > 0$*.*

Proof. We prove the case where the reduction R is quantum. The classical case can be done similarly. Let Π be the promise problem in the statement. Let $\mathcal{F}$ denote Protocol 24 that returns the two circuits in Lines 3 and 4, and h be its advice as follows: $h := (K_a, T_a, p, b_Y, b_N)$. The construction of the non-uniform EFI is the following:

- $\mathsf{EFI}_h(1^{\lambda}, b)$: Sample $y \sim \mathcal{U}_{T_a}$. Compute $(\widehat{C}_0, \widehat{C}_1) \leftarrow \mathcal{F}(y)$. Return the state $\widehat{C}_b|\mathbf{0}\rangle$.

Note that T_a has only YES instances.

The two output states are statistically far. By Theorem 8.5.1, the pair of circuits $(\widehat{C}_0, \widehat{C}_1) \leftarrow \mathcal{F}(y)$ is a $\mathrm{QSD}_{1-2\mu,\delta+\gamma}$ instance. Since $y \in \Pi_Y$, then $\|\widehat{C}_0|\mathbf{0}\rangle - \widehat{C}_1|\mathbf{0}\rangle\|_1 \geq 1 - 2\mu(\lambda)$. This concludes the statistical distinguishability.

On the computational indistinguishability, we will argue by contradiction. Assume there exists an adversary $\mathcal{A}$ that distinguishes the EFI states $\widehat{C}_b|\mathbf{0}\rangle$ with advantage $\nu(\lambda)$ that is to be determined later. Let us consider an algorithm $\mathcal{B}$ trying to solve Π as follows: given an instance $z \in \{0,1\}^{\lambda}$, it first computes $(C'_0, C'_1) \leftarrow \mathcal{F}(z)$, then it samples a uniform coin $b \leftarrow \{0,1\}$ and relays $C'_b|\mathbf{0}\rangle$ to the distinguisher $\mathcal{A}$. Finally, $\mathcal{B}$ will return 1 if $\mathcal{A}$ returns b, and 0 otherwise.

Case $z \in \Pi_Y$: Suppose that z has been sampled from $\mathcal{U}_{T_a}$. Then, the (mixed) state $C_b'|\mathbf{0}\rangle$ that we deliver to the adversary $\mathcal{A}$ would be identical to the EFI state $\widehat{C}_b|\mathbf{0}\rangle$. Therefore, from the $\nu(\lambda)$-distinguishability of EFI states for $\mathcal{A}$, we would have

$$\Pr(\mathcal{B}(z)=1) = \Pr(\mathcal{A}(P_b|0\rangle)=b) \geq \frac{1}{2}+\frac{\nu(\lambda)}{2} \ . \tag{8.68}$$

We know that z does not necessarily follow the distribution $\mathcal{U}_{T_a}$. However, one can argue that $\widehat{C}_b$ is not far from C_b' by leveraging the disguising lemma. We have that

$$\|\widehat{C}_0 \otimes \widehat{C}_1|\mathbf{0},\mathbf{0}\rangle - C_0' \otimes C_1'|\mathbf{0},\mathbf{0}\rangle\|_1 \tag{8.69}$$

$$\leq \|\widehat{C}_1|\mathbf{0}\rangle - C_1'|\mathbf{0}\rangle\|_1 \tag{8.70}$$

$$\leq \mathbb{E}_{a\sim\mathcal{U}_{[s]},\pi\sim\mathfrak{S}_m}\left[\left\|R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p},\mathcal{U}_{T_a}^{\otimes p}\right)\right)-R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p},y,\mathcal{U}_{T_a}^{\otimes p-1}\right)\right)\right\|_1\right] \tag{8.71}$$

$$\leq \delta(\alpha(\lambda)) + \frac{2(m+1)}{d+1} + \varepsilon \tag{8.72}$$

$$\leq \delta(\alpha(\lambda)) + \gamma \ , \tag{8.73}$$

where we used the fact that $\widehat{C}_0 = C_0'$, properties of the trace distance, and Lemma 8.3.3. Using the fact that the trace distance is decreasing under partial trace, for any $b \in \{0,1\}$, we obtain

$$\|\widehat{C}_b|\mathbf{0}\rangle - C_b'|\mathbf{0}\rangle\|_1 \leq \delta(\alpha(\lambda)) + \gamma \ . \tag{8.74}$$

The adversary $\mathcal{A}$ can thus distinguish the C_b''s with probability

$$\Pr(\mathcal{B}(z)=1) = \Pr\big(\mathcal{A}(C_b'|\mathbf{0}\rangle)=b\big) \tag{8.75}$$

$$= \frac{1}{2}+\frac{1}{2}\left|\Pr_{x\leftarrow C_0'}(\mathcal{A}(x)=1) - \Pr_{x\leftarrow C_1'}(\mathcal{A}(x)=1)\right| \tag{8.76}$$

$$\begin{aligned}\geq \frac{1}{2}+\frac{1}{2}\Bigg(&\left|\Pr_{x\leftarrow \widehat{C}_0}(\mathcal{A}(x)=1) - \Pr_{x\leftarrow \widehat{C}_1}(\mathcal{A}(x)=1)\right| \\ &-\left|\Pr_{x\leftarrow \widehat{C}_0}(\mathcal{A}(x)=1) - \Pr_{x\leftarrow C_0'}(\mathcal{A}(x)=1)\right| \\ &-\left|\Pr_{x\leftarrow \widehat{C}_1}(\mathcal{A}(x)=1) - \Pr_{x\leftarrow C_1'}(\mathcal{A}(x)=1)\right|\Bigg)\end{aligned} \tag{8.77}$$

$$\geq \frac{1}{2}+\frac{\nu(\lambda)}{2} - \delta(\alpha(\lambda)) - \gamma \ . \tag{8.78}$$

Case $z \in \Pi_N$: By Theorem 8.5.1, the evaluations of the two circuits (C_0', C_1') from the reduction $\mathcal{F}(z)$ are close in trace distance, namely,

$$\|C_0'|\mathbf{0}\rangle - C_1'|\mathbf{0}\rangle\|_1 \leq \delta(\alpha(\lambda)) + \gamma \ . \tag{8.79}$$

Recall that the trace distance provides the maximum distinguishability advantage for *any* distinguisher, including $\mathcal{A}$, therefore

$$\begin{aligned}\Pr(\mathcal{B}(z) = 1) = \Pr\big(\mathcal{A}(C_b'|\mathbf{0}\rangle) = b\big) &\leq \frac{1}{2}(1 + \|C_0'|\mathbf{0}\rangle - C_1'|\mathbf{0}\rangle\|_1) &(8.80)\\ &\leq \frac{1}{2}(1 + \delta(\alpha(\lambda)) + \gamma) \ . &(8.81)\end{aligned}$$

Conclusion: We need one more algorithm that will leverage the capacity of $\mathcal{B}$ to decide Π. Let $k \in \mathbb{N}$, and $\mathcal{C}$ be an algorithm that on instance $z \in \{0,1\}^\lambda$, runs $\mathcal{B}(z)$ for k times independently. Let $b_1, \dots, b_k$ be k corresponding independent outputs of $\mathcal{B}(z)$. Then $\mathcal{C}$ returns:

$$\begin{cases} 0 & \text{if } \left|\frac{1}{k}\sum_i b_i - \frac{1}{2}\right| \geq \tau(\lambda), \\ 1 & \text{otherwise,} \end{cases} \tag{8.82}$$

where $\tau(\lambda)$ is chosen such that

$$\tau(\lambda) := \frac{\nu(\lambda)}{4} - \frac{3(\delta(\lambda) + \gamma)}{4} \ . \tag{8.83}$$

Then, we have

$$\begin{aligned}&\Pr(\mathcal{C}(z) = 0 | z \in \Pi_Y) &(8.84)\\ &= \Pr\left(\left|\frac{1}{k}\sum_i b_i - \frac{1}{2}\right| \geq \tau(\lambda) \;\middle|\; b_1, \dots, b_k \leftarrow \mathcal{B}(z), z \in \Pi_Y\right) &(8.85)\\ &\geq \Pr\left(\frac{1}{k}\sum_i b_i \geq \frac{1}{2} + \tau(\lambda) \;\middle|\; b_1, \dots, b_k \leftarrow \mathcal{B}(z), z \in \Pi_Y\right) &(8.86)\\ &\geq \Pr\left(\frac{1}{k}\left(\sum_i b_i - \mathbb{E}(\mathcal{B}_i(z))\right) \geq -\tau(\lambda) \;\middle|\; b_1, \dots, b_k \leftarrow \mathcal{B}(z), z \in \Pi_Y\right) &(8.87)\\ &\geq 1 - \exp\big(-2k\tau(\lambda)^2\big) \ , &(8.88)\end{aligned}$$

where we used $\mathbb{E}(\mathcal{B}_i(z)) - \tau(\lambda) \geq \frac{1}{2} + \tau(\lambda)$ for $z \in \Pi_Y$ by equation (8.75) in the second inequality, and Hoeffding's lemma in the last inequality. On the other hand, we have

$$\Pr(\mathcal{C}(z) = 1 | z \in \Pi_N) \tag{8.89}$$

$$= \Pr\left(\left|\frac{1}{k}\sum_i b_i - \frac{1}{2}\right| < \tau(\lambda) \;\middle|\; b_1, \ldots, b_k \leftarrow \mathcal{B}(z), z \in \Pi_N\right) \tag{8.90}$$

$$= \Pr\left(\left|\frac{1}{k}\sum_i b_i - \frac{1}{k}\sum_i \mathbb{E}(\mathcal{B}_i(z))\right| < \tau(\lambda) \;\middle|\; b_1, \ldots, b_k \leftarrow \mathcal{B}(z), z \in \Pi_N\right) \tag{8.91}$$

$$\geq 1 - \exp\bigl(-2k\tau(\lambda)^2\bigr) , \tag{8.92}$$

where we once again used Hoeffding's lemma and equation (8.80). For $k := 1/\tau(\lambda)^2$, any sufficiently large $\lambda \in \mathbb{N}$, and any $z \in (\Pi_Y \cup \Pi_N) \cap \{0,1\}^\lambda$, it holds that

$$\Pr(\mathcal{C}(z) = \chi_\Pi(z)) \geq 1 - \exp\bigl(-2k\tau(\lambda)^2\bigr) \geq \frac{2}{3} , \tag{8.93}$$

This breaks the worst-case hardness of Π.

Since $\theta_{\mathsf{efi}} := (1-2\mu(\lambda)) - 3(\delta(\alpha(\lambda))+\gamma)$, we can set $\nu(\lambda) := (1-2\mu(\alpha)) - \theta_{\mathsf{efi}}/2$, and the number of repetitions in the last step becomes

$$\frac{1}{\tau(\lambda)^2} = \frac{4^2}{(\nu(\lambda) - 3(\delta(\alpha(\lambda)) + \gamma))^2} = \frac{4^3}{\theta_{\mathsf{efi}}^2} . \tag{8.94}$$

Runtime: We compute the runtime of EFI as follows. It first samples $2m$ instances from $\mathcal{U}_{K_a}$ (or $\mathcal{U}_{T_a}$), applies the permutations π twice to each half of the samplings, and computes R on each half. One single sampling from $\mathcal{U}_{K_a}$ (or $\mathcal{U}_{T_a}$) takes time $O(d\lambda)$, where $d \leq (m+1)/\gamma$ is the size of $\mathcal{U}_{K_a}$ and λ is the size of each element in $\mathcal{U}_{K_a}$. The permutations can be applied in time $O(m)$. Therefore, the total runtime of EFI is $O(T + m^2\lambda/\gamma)$.

Note that $\mathcal{C}$ runs $\mathcal{B}$ for $O(1/{\theta_{\mathsf{efi}}}^2)$ times. Each execution of $\mathcal{B}$ evaluates $\widehat{C}_b$, queries $\mathcal{A}$, and performs an equality check. All of this takes $O((T + m^2\lambda/\gamma)/{\theta_{\mathsf{efi}}}^2)$ with $O(1/{\theta_{\mathsf{efi}}}^2)$ queries to $\mathcal{A}$.

□

8.6.2 One-way state generators from mildly-lossy problems

In the next theorem, we discuss the adaptation to the quantum setting, when λ is relatively large.

Theorem 8.6.2. *Let Π be a $(T, \mu, f^m, \alpha, \gamma)$-mildly-lossy problem with a pure-outcome reduction. Also assume that $\theta_{\mathsf{ows}} := 1-(\delta(\alpha(\lambda))+\gamma+4\sqrt{2\mu(\lambda)}) > 0$ and $\tau_{\mathsf{ows}} := 1 - 2\mu(\lambda) - (\delta(\alpha(\lambda)) + \gamma) > 0$, with δ as in Definition 8.4.7. Then there exists an algorithm $\mathsf{G} = (\mathsf{StateGen}, \mathsf{Ver})$ such that $\mathsf{StateGen}$ runs in time $O(T + m^2\lambda\gamma^{-1})$ and an oracle algorithm $\mathcal{C}$, such that for every algorithm $\mathcal{A}$ one and only one of the following statements holds:*

i. $\mathcal{C}^{\mathcal{A}}$ solves $\Pi \cap \{0,1\}^n$ in time $O((T + m^2\lambda\gamma^{-1} + \tau_{\mathsf{ows}}^{-2})\theta_{\mathsf{ows}}^{-2})$ with $O(\theta_{\mathsf{ows}}^{-2})$ classical queries to $\mathcal{A}$,

ii. G is a $(1 - \theta_{\mathsf{ows}}/4)$-OWSG for $\mathcal{A}$.

Proof. Sample $z \sim \mathcal{U}_{K_a}$ and apply Protocol 24 up to Line 4 on input z to obtain the two circuits (C_0^*, C_1^*). Note that the two circuits are mixed; classical randomness is used to sample $\widetilde{x}$ but the algorithm R is a pure quantum circuit. Let κ be the size of the randomness of these circuits. For any $r \in \{0,1\}^\kappa$ and $b \in \{0,1\}$, let $C_b^*|r, \mathbf{0}\rangle$ be the pure state obtained by sampling $\widetilde{x}$ using r and applying R to $\pi(\widetilde{x})$ and a possibly ancilla $|\mathbf{0}\rangle$ with an appropriate size. We show that G, defined as follows:

- $\mathsf{StateGen}(r, b)$: output $C_b^*|r, \mathbf{0}\rangle$.
- $\mathsf{Ver}((r, b), \rho)$: If $\|C_b^*|r, \mathbf{0}\rangle - \rho\|_1 \leq \delta(\alpha(\lambda)) + \gamma$ output $\top$, otherwise output $\perp$.

is a $(\theta_{\mathsf{ows}}/2)$-weak one-way state generator.

Assume that there exists an adversary $\mathcal{A}$ that breaks the scheme above with probability more than $1 - \theta_{\mathsf{ows}}/4$. We use $\mathcal{A}$ to construct an algorithm for Π. Consider the following oracle algorithm $\mathcal{B}^{\mathcal{A}}$:

- $\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y)$: computes $(\widehat{C}_0, \widehat{C}_1(y))$ as in Protocol 24 up to Line 4 on input y. Samples a uniform $r \in \{0,1\}^\kappa$ and a uniform $b \in \{0,1\}$, and computes $\rho := \widehat{C}_b|r, \mathbf{0}\rangle$. Runs the adversary $(r', b') \leftarrow \mathcal{A}(\rho)$, and computes $\rho' = \widehat{C}_{b'}|r', \mathbf{0}\rangle$. If $\|\rho - \rho'\|_1 \leq \delta + \gamma$ it outputs $\top$, otherwise it outputs $\perp$.

We compute the advantage of $\mathcal{B}$ in distinguishing between YES and NO instances of Π by analysing the probability $\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = \top\Big)$.

Case $y \in \Pi_Y$: We show that for every $y \in \Pi_Y$, we have:

$$\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = \top\Big) \leq \frac{1}{2} + 2\sqrt{2\mu(\lambda)}\ . \tag{8.95}$$

Instead of proving the inequality directly for the circuits $(\widehat{C}_0, \widehat{C}_1)$, we will show it for two similar circuits $(\widetilde{C}_0, \widetilde{C}_1)$ with disjoint images. Let $\widehat{\rho}_0$ and $\widehat{\rho}_1$ be respectively the mixed states $\widehat{C}_0|r, \mathbf{0}\rangle$ and $\widehat{C}_1|r, \mathbf{0}\rangle$ when r follows the uniform distribution. For any POVM $\mathcal{M} = \{M_i\}_i$, let us define by $A_{\mathcal{M}}$ the following set:

$$A_{\mathcal{M}} := \{i \mid \operatorname{tr}(M_i\widehat{\rho}_0) \geq \operatorname{tr}(M_i\widehat{\rho}_1)\} . \tag{8.96}$$

In Theorem 8.5.1, we showed that for every $y \in \Pi_Y$, the statistical distance between $\widehat{\rho}_0$ and $\widehat{\rho}_1$ is at least $1 - 2\mu(\lambda)$. Moreover, we can rewrite the trace distance in terms of the POVMs as

$$\begin{aligned}
&\|\widehat{\rho}_0 - \widehat{\rho}_1\|_1 && (8.97)\\
&= \max_{\{M_i\}_i} \frac{1}{2} \sum_i |\operatorname{tr}(M_i\widehat{\rho}_0) - \operatorname{tr}(M_i\widehat{\rho}_1)| && (8.98)\\
&= \max_{\{M_i\}_i} \frac{1}{2} \left[\sum_{i \in A_{\mathcal{M}}} (\operatorname{tr}(M_i\widehat{\rho}_0) - \operatorname{tr}(M_i\widehat{\rho}_1)) + \sum_{i \in A^c_{\mathcal{M}}} (\operatorname{tr}(M_i\widehat{\rho}_1) - \operatorname{tr}(M_i\widehat{\rho}_0)) \right] && (8.99)\\
&= \max_{\{M_i\}_i} \left\{ \sum_{i \in A_{\mathcal{M}}} \operatorname{tr}(M_i\widehat{\rho}_0) + \sum_{i \in A^c_{\mathcal{M}}} \operatorname{tr}(M_i\widehat{\rho}_1) - 1 \right\} . && (8.100)
\end{aligned}$$

It follows that there exists a particular POVM $\mathcal{M}$, such that if we define the projections of $\widehat{C}_0$ and $\widehat{C}_1$ onto $A_{\mathcal{M}}$ and $A^c_{\mathcal{M}}$ by $\widetilde{C}_0$ and $\widetilde{C}_1$ respectively, i.e.,

$$\widetilde{C}_0 = \sum_{i \in A_{\mathcal{M}}} M_i\widehat{C}_0, \quad \text{and} \quad \widetilde{C}_1 = \sum_{i \in A^c_{\mathcal{M}}} M_i\widehat{C}_1 ,$$

we have $\operatorname{tr}(\widetilde{\rho}_0) + \operatorname{tr}(\widetilde{\rho}_1) \geq 2 - 2\mu(\lambda)$, where $\widetilde{\rho}_b$ is the mixed state $\widetilde{C}_b|r, \mathbf{0}\rangle$ and r is uniform. Therefore

$$(\operatorname{tr}(\widetilde{\rho}_0) \geq 1 - \mu(\lambda)) \wedge (\operatorname{tr}(\widetilde{\rho}_1) \geq 1 - 2\mu(\lambda)) , \tag{8.101}$$

or,

$$(\operatorname{tr}(\widetilde{\rho}_0) \geq 1 - 2\mu(\lambda)) \wedge (\operatorname{tr}(\widetilde{\rho}_1) \geq 1 - \mu(\lambda)) . \tag{8.102}$$

By the Gentle Measurement Lemma (Lemma 8.2.6), for either of the cases above, we have

$$\|\widehat{\rho}_0 - \widetilde{\rho}_0\|_1 \leq \sqrt{2\mu(\lambda)} , \quad \text{and} \quad \|\widehat{\rho}_1 - \widetilde{\rho}_1\|_1 \leq \sqrt{2\mu(\lambda)} . \tag{8.103}$$

Pretend that $\mathcal{A}$ also tried to distinguish between for $\widehat{C}_b$ and $\widetilde{C}_b$ for $b \in \{0,1\}$, and consider the following sequence of games that modifies $\mathcal{B}^{\mathcal{A}}$.

i. $\mathcal{G}_1$: In this game $\mathcal{B}$ behaves originally as above.

ii. $\mathcal{G}_2$: In this game $\mathcal{B}$ replaces $\widehat{C}_0$ with $\widetilde{C}_0$. Note that $\mathcal{A}$ can distinguish this modification with probability at most $\sqrt{2\mu(\lambda)}$ according to eq. (8.103). It follows that

$$\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = \top\Big) \tag{8.104}$$

$$\leq \Big|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = \top\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = \top\Big)\Big| \tag{8.105}$$

$$+ \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = \top\Big) \tag{8.106}$$

$$\leq \sqrt{2\mu(\lambda)} + \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = \top\Big)\,. \tag{8.107}$$

iii. $\mathcal{G}_3$: In this game, $\mathcal{B}$ replaces $\widehat{C}_1$ with $\widetilde{C}_1$. Note that $\mathcal{A}$ can identify this modification with probability at most $\sqrt{2\mu(\lambda)}$. We obtain

$$\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = \top\Big) \tag{8.108}$$

$$\begin{aligned}\leq &\Big|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = \top\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = \top\Big)\Big| \\ &+ \Big|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = \top\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = \top\Big)\Big| \\ &+ \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = \top\Big)\end{aligned} \tag{8.109}$$

$$\leq 2\sqrt{2\mu(\lambda)} + \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = \top\Big) \tag{8.110}$$

Now, note that the projections onto the supports of $\widetilde{C}_0$ and $\widetilde{C}_1$ are orthogonal to each other. Therefore, the adversary never succeeds when the bit b (chosen by $\mathcal{B}$) is equal to 1; there exists no r' such that $\|\widetilde{C}_0|r, |\mathbf{0}\rangle\rangle - \widetilde{C}_1|r', \mathbf{0}\rangle\|_1 \leq \delta(\alpha(\lambda)) + \gamma$. So

$$\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = \top\Big) \tag{8.111}$$

$$= \frac{1}{2}\Big(\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = \top | b = 0\Big) + \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = \top | b = 1\Big)\Big) \leq \frac{1}{2}\,. \tag{8.112}$$

Case $y \in \Pi_N$: By Lemma 8.3.3, the trace distance of the outcomes of $\widehat{C}_1$ and C_1^* is at most $\delta(\alpha(\lambda)) + \gamma$. Moreover, $\widehat{C}_0$ is exactly the same as C_0^*. Therefore, if the bit b chosen by $\mathcal{B}$ is equal to 0, then $\mathcal{A}$ succeeds with probability at least $1 - \theta_{\mathsf{ows}}/4$, and if $b = 1$, it succeeds with probability $1 - \theta_{\mathsf{ows}}/4 - (\delta(\alpha(\lambda)) + \gamma)$. In total, we obtain

$$\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = \top\Big) \geq \frac{1}{2}(1 - \frac{\theta_{\mathsf{ows}}}{4}) + \frac{1}{2}(1 - \frac{\theta_{\mathsf{ows}}}{4} - (\delta(\alpha(\lambda)) + \gamma)) \tag{8.113}$$

$$= 1 - \frac{\theta_{\mathsf{ows}}}{4} - \frac{(\delta(\alpha(\lambda)) + \gamma)}{2} . \tag{8.114}$$

Conclusion: We showed that the quantity of $\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = \top\Big)$ diverges for YES and NO instances of y. For our choice of parameters, we have

$$1 - \frac{\theta_{\mathsf{ows}}}{4} - \frac{(\delta(\alpha(\lambda)) + \gamma)}{2} - \left(\frac{1}{2} + 2\sqrt{2\mu(\lambda)}\right) \tag{8.115}$$

$$= \frac{1 - (\delta(\alpha(\lambda)) + \gamma + 4\sqrt{2\mu(\lambda)})}{2} - \frac{\theta_{\mathsf{ows}}}{4} \tag{8.116}$$

$$= \frac{\theta_{\mathsf{ows}}}{4} . \tag{8.117}$$

Let $\mathcal{C}$ be an algorithm that runs $\mathcal{B}$ for $O(1/\theta_{\mathsf{ows}}^2)$ many times, and approximates the quantity above within error less than $\theta_{\mathsf{ows}}/4$. If this value is more than $1 - \theta_{\mathsf{ows}}/4 - (\delta(\alpha(\lambda)) + \gamma)/2$, then y must be a NO instance, otherwise it is a YES instance. Therefore, we finally obtain a algorithm that solves Π. Note that $\mathcal{B}$ verifies whether $\|\widehat{C}_b|r, \mathbf{0}\rangle - \widehat{C}_{b'}|r', \mathbf{0}\rangle\|_1$ is smaller than $\delta(\alpha(\lambda)) + \gamma$. Since the reduction R is pure and r, r' are fixed, these states are pure, therefore $\mathcal{B}$ can perform a swap test Lemma 8.2.7 for $O(1/\tau_{\mathsf{ows}}^2)$ number of times on them to approximate their ℓ_1 distance. □

8.7 Worst-case to average-case reductions

In Section 8.4 we introduced mildly-lossy problems, promise problems that admit reductions that *lose* some information about the input, and in Section 8.6 we constructed cryptography primitives from these.

In this section we analyse the mild-lossiness of worst-case to average-case reductions. Since we discuss mild-lossiness of such reductions, as motivated in Section 8.4, we focus on worst-case to average-case *f-distinguisher* reductions (Definition 8.4.1). In Definition 8.7.1, we put forward the definition of *worst-case to average-case f-distinguisher reduction* which can be viewed as a generalization of worst-case to average-case reductions in the sense that (i) the reduction is oblivious to the target average-case problem (inherited from being f-distinguisher), and (ii) the reduction maps inputs to a distribution that is *not* necessarily efficiently samplable. The latter does not impose any issues in our setting, since we are only discussing mild-lossiness of the reductions, and not the hardness of the problems. We then prove, in Theorem 8.7.2, that such reductions are lossy and specify the mild-lossiness parameters. Combined with Theorems 8.6.1 and 8.6.2, this would yield in Corollaries 8.7.3 and 8.7.4 that worst-case to average-case reductions can be used to build EFIs and OWSGs.

Definition 8.7.1 (Worst-Case to Average-Case f-Distinguisher Reduction)**.** *Let Π be a promise problem, $\lambda \in \mathbb{N}$, and $d \in [0,1]$. We say that a reduction R is a (T, μ, f^m, d)-worst-case to average-case (WC-DIST) reduction for Π if*

i. R is a (μ, f^m)-distinguisher reduction for Π (per Definition 8.4.1),

ii. for all $x \in \Pi \cap \{0,1\}^\lambda$, $R(x)$ runs in time $T(\lambda)$,

iii. there exists a distribution $D = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ over $\{0,1\}^*$, such that

$$\forall (x_1, \cdots, x_m) \in (\Pi \cap \{0,1\}^\lambda)^m : \frac{1}{2}\|R(x_1, \cdots, x_m) - D\|_1 \leq d\,. \quad (8.118)$$

The upper bound d is called the distance of the reduction.

If there exist two distributions D_Y and D_N over $\{0,1\}^*$ such that for inputs $x \in \Pi_Y$ the distribution D_Y approximates $R(x)$ up to error d, and for inputs $x \in \Pi_N$ the distribution D_N approximates $R(x)$ up to error d, we say that the reduction R is a (T, μ, f^m, d)-worst-case to average-case splitting-reduction for Π.

We now show that a worst-case to average-case quantum reduction implies mild-lossiness, and therefore OWSGs and EFIs. However, since a quantum reverse Pinsker inequality is not known in its most general form, we include here two independent assumptions on the quantum worst-case to average-case reductions that imply quantum cryptography.

Theorem 8.7.2 (Mild-Lossiness of WC-DIST f-Distinguisher Quantum Reductions)**.** *Let $\Pi = \Pi_Y \cup \Pi_N$ for two disjoint sets $\Pi_Y, \Pi_N \subset \{0,1\}^*$. If there exists a (T, μ, f^m, d)-WC-DIST quantum reduction R for Π, such that f is a non-constant permutation-invariant function, then for any $\gamma > 0$, we have*

i. If the minimum eigenvalue of the reduction is uniformly bounded from below for every pairwise independent $2^9 m\lambda/\gamma^3$-uniform distribution X, i.e. there exists a constant $\beta(\lambda) > 0$ such that $\lambda_{\min}(R(X)) > \beta(\lambda)$, then Π is $(T, \mu, f^m, \alpha, \gamma)$-mildly-lossy, where

$$\alpha(\lambda) = (\beta(\lambda) + 2d) \log\left(1 + \frac{2d}{\beta(\lambda)}\right). \tag{8.119}$$

ii. If instead the dimension of the image space is upper bounded for every pairwise independent $2^9 m\lambda/\gamma^3$-uniform distribution X, i.e. there exists a constant $d_R(\lambda) \in \mathbb{N}$ such that $\dim(\mathrm{Im}(R(X))) \leq d_R(\lambda)$, then Π is $(T, \mu, f^m, \alpha, \gamma)$-mildly-lossy, where

$$\alpha(\lambda) = 4d \log d_R(\lambda) + h(2d). \tag{8.120}$$

Proof. We will prove it separately for both cases.

Case 1: $\lambda_{\min}(R(X)) > \beta(\lambda)$. Let us denote by $\rho_{X,R(X)}$ (or simply by ρ) the joint system of the classical-quantum state after the reduction R is applied to a pairwise independent $2^9 m\lambda/\gamma^3$-uniform distribution X_λ over λ-bit strings, where we drop the subscript λ for simplicity, see eq. (8.5). We denote the subsystems of $\rho_{X,R(X)}$ by ρ_X and $\rho_{R(X)}$. Note that since $\rho_{X,R(X)}$ is a classical-quantum system, so is $\rho_X \otimes \rho_{R(X)}$. We can rewrite the mutual information in terms of the relative entropy, which by eq. (8.9) for classical-quantum systems takes a simple form

$$\begin{aligned} I(X; R(X))_\rho &= D(\rho_{X,R(X)} \parallel \rho_X \otimes \rho_{R(X)}) & (8.121)\\ &= \sum_x \Pr(X = x) D_{KL}(\rho_{R(X)|X=x} \parallel \rho_{R(X)}) \,. & (8.122) \end{aligned}$$

Note that we drop the classical term from the previous equation because both states have the same classical distribution. By Lemma 8.2.3, if the minimum eigenvalue of the reduction is uniformly bounded from below by a constant $\beta(\lambda)$, i.e. $\lambda_{\min}(R(X)) > \beta(\lambda)$, then we have a reverse Pinsker-like inequality

$$\begin{aligned} D(\rho_{R(X)|X=x}||\rho_{R(X)}) \\ &\leq \left(\beta(\lambda) + \frac{1}{2}\|\rho_{R(X)|X=x} - \rho_{R(X)}\|_1\right) \\ &\quad \cdot \log\left(1 + \frac{1}{2\beta(\lambda)}\|\rho_{R(X)|X=x} - \rho_{R(X)}\|_1\right) . \end{aligned} \tag{8.123}$$

Finally, note that since R is a (T, μ, f^m, d)-WC-DIST reduction, there exists a distribution D_λ such that for any $x \in \Pi \cap \{0,1\}^\lambda$, it holds that $\frac{1}{2}\|\rho_{R(X)|X=x} - D_\lambda\|_1 \leq d$, thus $\frac{1}{2}\|\rho_{R(X_\lambda)} - D_\lambda\|_1 \leq d$. By the triangle inequality, we have $\frac{1}{2}\|\rho_{R(X)|X=x} - \rho_{R(X)}\|_1 \leq 2d$. We conclude that

$$I(X; R(X))_\rho \leq \sum_x \Pr(X = x)(\beta(\lambda) + 2d) \log\left(1 + \frac{2d}{\beta(\lambda)}\right) \tag{8.124}$$

$$= (\beta(\lambda) + 2d) \log\left(1 + \frac{2d}{\beta(\lambda)}\right) . \tag{8.125}$$

Case 2: $\dim(\mathrm{Im}(R(X))) \leq d_R(\lambda)$**.** We can find an alternative bound using the quantum conditional entropy. Let us denote by ω the product state $\omega_{X,R(X)} := \rho_X \otimes \rho_{R(X)}$, since the mutual information between subsystems of product states are zero, we have

$$I(X; R(X))_\rho = |I(X; R(X))_\rho - I(X; R(X))_\omega| \tag{8.126}$$

$$= |S(\rho_X) - S(X|R(X))_\rho - S(\omega_X) + S(X|R(X))_\omega| \tag{8.127}$$

$$= |S(X|R(X))_\rho - S(X|R(X))_\omega| \tag{8.128}$$

$$\leq 2\mathrm{TD}(\rho, \omega) \log \dim(H_A) + h(\mathrm{TD}(\rho, \omega)) , \tag{8.129}$$

where in the last inequality we used Lemma 8.2.5. We can bound the trace distance between ρ and ω by the worst-case indistinguishability of the reduction R. Indeed, note that ρ and ω are classical-quantum states with the same classical distribution, thus

$$\|\rho_{X,R(X)} - \rho_X \otimes \rho_{R(X)}\|_1 = \sum_x \Pr(X = x)\|\rho_{R(X)|X=x} - \rho_{R(X)}\|_1 \tag{8.130}$$

$$\leq \sum_x \Pr(X = x)2d = 2d . \tag{8.131}$$

Since the binary entropy function is increasing on $[0, 1/2]$, we can conclude that for $d < 1/4$,

$$I(X; R(X))_\rho \leq 4d \log d_R(\lambda) + h(2d) \,. \tag{8.132}$$

□

The following corollaries stating conditions for the existence of OWSG are a direct result of combining Theorems 8.6.2 and 8.7.2, we split the two conditions on the quantum reduction for clarity.

Corollary 8.7.3. *Let Π be a promise problem, and assume that there exists a (T, μ, f^m, d)-WC-DIST reduction for Π. Let $\beta(\lambda) > 0$ be a uniform bound for the minimum eigenvalue of the reduction $\lambda_{\min}(R(X)) > \beta(\lambda)$ for every pairwise independent $2^9 m\lambda/\gamma^3$-uniform distribution X. Let $\theta_{\mathsf{ows}} := 1 - (\delta(\alpha(\lambda)) + \gamma + 4\sqrt{2\mu(\lambda)}) > 0$ and $\tau_{\mathsf{ows}} := 1 - 2\mu(\lambda) - (\delta(\alpha(\lambda)) + \gamma) > 0$, where $\gamma > 0$, $\alpha(\lambda) = (\beta(\lambda) + 2d)\log(1 + 2d/\beta(\lambda))$, and δ is the function defined in Definition 8.4.7. Then there exists an algorithm $\mathsf{G} = (\mathsf{StateGen}, \mathsf{Ver})$ such that $\mathsf{StateGen}$ runs in time $O(T + m^2\lambda\gamma^{-1})$ and an oracle algorithm $\mathcal{C}$, such that for every algorithm $\mathcal{A}$ one and only one of the following statements holds:*

i. $\mathcal{C}^{\mathcal{A}}$ solves $\Pi \cap \{0,1\}^\lambda$ in time $O((T + m^2\lambda\gamma^{-1} + \tau_{\mathsf{ows}}^{-2})\theta_{\mathsf{ows}}^{-2})$ with $O(\theta_{\mathsf{ows}}^{-2})$ classical queries to $\mathcal{A}$,

ii. G is a $(1 - \theta_{\mathsf{ows}}/4)$-OWSG for $\mathcal{A}$.

Corollary 8.7.4. *Let Π be a promise problem, and assume that there exists a (T, μ, f^m, d)-WC-DIST reduction for Π, with $d < 1/4$. Let $d_R(\lambda) \in \mathbb{N}$ be such that $\dim(\mathrm{Im}(R(X))) \leq d_R(\lambda)$ for every pairwise independent $2^9 m\lambda/\gamma^3$-uniform distribution X. Let $\theta_{\mathsf{ows}} := 1 - (\delta(\alpha(\lambda)) + \gamma + 4\sqrt{2\mu(\lambda)}) > 0$ and $\tau_{\mathsf{ows}} := 1 - 2\mu(\lambda) - (\delta(\alpha(\lambda)) + \gamma) > 0$, where $\gamma > 0$, $\alpha(\lambda) = 4d \log d_R(\lambda) + h(2d)$, and δ is the function defined in Definition 8.4.7. Then there exists an algorithm $\mathsf{G} = (\mathsf{StateGen}, \mathsf{Ver})$ such that $\mathsf{StateGen}$ runs in time $O(T + m^2\lambda\gamma^{-1})$ and an oracle algorithm $\mathcal{C}$, such that for every algorithm $\mathcal{A}$ one and only one of the following statements holds:*

i. $\mathcal{C}^{\mathcal{A}}$ solves $\Pi \cap \{0,1\}^\lambda$ in time $O((T + m^2\lambda\gamma^{-1} + \tau_{\mathsf{ows}}^{-2})\theta_{\mathsf{ows}}^{-2})$ with $O(\theta_{\mathsf{ows}}^{-2})$ classical queries to $\mathcal{A}$,

ii. G is a $(1 - \theta_{\mathsf{ows}}/4)$-OWSG for $\mathcal{A}$.

Remark 8.7.1. *We can also instantiate Theorem 8.7.2 with the construction of EFIs in Theorem 8.6.1 to obtain* $(1-2\mu(\lambda), 1-2\mu(\lambda)-\theta_{\mathsf{efi}}/2)$*-EFIs from WC-DIST* f*-Distinguisher Quantum Reductions with the same two possible conditions on the parameter* $\alpha(\lambda)$ *from* $\theta_{\mathsf{efi}} := (1-2\mu(\lambda)) - 3(\delta(\alpha(\lambda)) + \gamma)$.

8.8 Conclusion

In this chapter we put forward the first construction of quantum cryptographic primitives (without auxiliary inputs) based on quantum reductions. We do so by generalizing Drucker's [Dru12] disguising distribution lemma for mildly-lossy reductions in Section 8.4 and by showing how to obtain zero-knowledgeness from these reductions in Section 8.5. Moreover, in order to support the belief in the existence of our reductions, we prove that worst-case to average-case reductions are mildly-lossy.

Note that although lossy reductions were introduced by Ball, Boyle, Degwekar, Deshpande, Rosen, Vaikuntanathan, and Vasudevan [BBD+20] they did not consider quantum reductions. In contrast to their work, we also show the existence of our primitives from the circuits obtained through zero-knowledgeness, which we believe makes them less opaque. Interestingly, they were able to construct *collision-resistant hash functions* (CRHF) from lossy reductions, which are believed to be harder than one-way functions. However, when trying to reproduce such a result in the quantum case, with a sort of computational variant of quantum fingerprinting [BCWdW01], we fall into the problem of efficient-verifiability of a "collision".

Given the diverse landscape of quantum cryptographic primitives, see figure 6.1, this thesis chapter only opens the door to many other research questions, to name a few:

- Given that quantum mildly-lossy reductions are much more relaxed than say compressing reductions, is there are any relevant problem (such as 3Sat or SVP) for which we can give a reduction?

- In this work we propose a reduction for EFIs and OWSGs, which notably have different lossiness requirements, it is therefore interesting to ask if either of these requirements imply pseudorandom state generators (PRSGs) or these require *even harder* complexity-theoretical assumptions.

- Although we were not able to prove the existence of collision-resistant *state* functions, or equivalently computational quantum fingerprints, it does not mean they cannot be defined. On the one hand, classically, CRHFs imply non-interactive commitments, which are separated from interactive ones. On the other hand, quantumly, non-interactive and interactive quantum commitments are equivalent. Therefore, we wonder what a relevant formulation of this quantum variant of CRHFs would be, and where it would sit in the hierarchy of quantum primitives.

Bibliography

[Aar04] Scott Aaronson. Limitations of quantum advice and one-way communication. In *19th IEEE Annual Conference on Computational Complexity*, pages 320–332. IEEE, 2004.

[ABKK23] Amit Agarwal, James Bartusek, Dakshita Khurana, and Nishant Kumar. A new framework for quantum oblivious transfer. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 363–394, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland.

[ABOE08] Dorit Aharonov, Michael Ben-Or, and Elad Eban. Interactive Proofs For Quantum Computations, 2008.

[ADSS17] Gorjan Alagic, Yfke Dulek, Christian Schaffner, and Florian Speelman. Quantum fully homomorphic encryption with verification. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 438–467, Hong Kong, China, December 3–7, 2017. Springer, Cham, Switzerland.

[AE11] Koenraad MR Audenaert and Jens Eisert. Continuity bounds on the quantum relative entropy—ii. *Journal of Mathematical Physics*, 52:112201, 2011.

[AGKL24] Prabhanjan Ananth, Aditya Gulati, Fatih Kaleoglu, and Yao-Ting Lin. Pseudorandom isometries. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part IV*, volume 14654 of *LNCS*, pages 226–254, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland.

[AGL24] Prabhanjan Ananth, Aditya Gulati, and Yao-Ting Lin. Cryptography in the common Haar state model: Feasibility results and separations. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part II*, volume 15365 of *LNCS*, pages 94–125, Milan, Italy, December 2–6, 2024. Springer, Cham, Switzerland.

[AGM18] Gorjan Alagic, Tommaso Gagliardoni, and Christian Majenz. Unforgeable quantum encryption. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 489–519, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Cham, Switzerland.

[AGM21] Gorjan Alagic, Tommaso Gagliardoni, and Christian Majenz. Can you sign a quantum state? *Quantum*, 5:603, 2021.

[AGQY22] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 237–265, Chicago, IL, USA, November 7–10, 2022. Springer, Cham, Switzerland.

[AK07] Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 115–128. IEEE, 2007.

[AL20] Prabhanjan Ananth and Rolando L. La Placa. Secure quantum extraction protocols. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 123–152, Durham, NC, USA, November 16–19, 2020. Springer, Cham, Switzerland.

[ALY24] Prabhanjan Ananth, Yao-Ting Lin, and Henry Yuen. Pseudorandom strings from pseudorandom quantum states. In Venkatesan Guruswami, editor, *ITCS 2024*, volume 287, pages 6:1–6:22, Berkeley, CA, USA, January 30 – February 2, 2024. LIPIcs.

[ANTV99] Andris Ambainis, Ashwin Nayak, Amnon Ta-Shma, and Umesh V. Vazirani. Dense quantum coding and a lower bound for 1-way quantum automata. In *31st ACM STOC*, pages 376–383, Atlanta, GA, USA, May 1–4, 1999. ACM Press.

[AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 208–236, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.

[BB21] Nir Bitansky and Zvika Brakerski. Classical binding for quantum commitments. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part I*, volume 13042 of *LNCS*, pages 273–298, Raleigh, NC, USA, November 8–11, 2021. Springer, Cham, Switzerland.

[BBC$^+$93] 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. *Physical review letters*, 70(13):1895–1899, 1993.

[BBD$^+$20] Marshall Ball, Elette Boyle, Akshay Degwekar, Apoorvaa Deshpande, Alon Rosen, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. Cryptography from information loss. In Thomas Vidick, editor, *ITCS 2020*, volume 151, pages 81:1–81:27, Seattle, WA, USA, January 12–14, 2020. LIPIcs.

[BBO$^+$24] Mohammed Barhoush, Amit Behera, Lior Ozer, Louis Salvail, and Or Sattath. Signatures from pseudorandom states via $\perp$-prfs, 2024.

[BCG$^+$02] Howard Barnum, Claude Crépeau, Daniel Gottesman, Adam Smith, and Alain Tapp. Authentication of quantum messages. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 449–458, 2002.

[BCG$^+$06] Michael Ben-Or, Claude Crépeau, Daniel Gottesman, Avinatan Hassidim, and Adam Smith. Secure multiparty quantum computation with (only) a strict honest majority. In *47th FOCS*, pages 249–260, Berkeley, CA, USA, October 21–24, 2006. IEEE Computer Society Press.

[BCG$^+$16] Harry Buhrman, Łukasz Czekaj, Andrzej Grudka, Michał Horodecki, Paweł Horodecki, Marcin Markiewicz, Florian Speelman, and Sergii Strelchuk. Quantum communication complexity advantage implies violation of a bell inequality. *Proceedings of the National Academy of Sciences*, 113(12):3191–3196, 2016.

[BCKM21] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 467–496, Virtual Event, August 16–20, 2021. Springer, Cham, Switzerland.

[BCN25] John Bostanci, Boyang Chen, and Barak Nehoran. Oracle separation between quantum commitments and quantum one-wayness. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 3–22, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland.

[BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In Yael Tauman Kalai, editor, *ITCS 2023*, volume 251, pages 24:1–24:21, Cambridge, MA, USA, January 10–13, 2023. LIPIcs.

[BCWdW01] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87:167902, Sep 2001.

[BD18] Zvika Brakerski and Nico Döttling. Two-message statistically sender-private OT from LWE. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 370–390, Panaji, India, November 11–14, 2018. Springer, Cham, Switzerland.

[BDRV19] Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Statistical difference beyond the polarizing regime. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 311–332, Nuremberg, Germany, December 1–5, 2019. Springer, Cham, Switzerland.

[BDS$^+$01] Charles H. Bennett, David P. DiVincenzo, Peter W. Shor, John A. Smolin, Barbara M. Terhal, and William K. Wootters. Remote state preparation. *Physical review letters*, 87(7):077902, 2001.

[BEM24] Samuel Bouaziz-Ermann and Garazi Muguruza. Quantum pseudorandomness cannot be shrunk in a black-box way. Cryptology ePrint Archive, Report 2024/291, 2024.

[BFV20] Adam Bouland, Bill Fefferman, and Umesh V. Vazirani. Computational pseudorandomness, the wormhole growth paradox, and constraints on the AdS/CFT duality (abstract). In Thomas Vidick, editor, *ITCS 2020*, volume 151, pages 63:1–63:2, Seattle, WA, USA, January 12–14, 2020. LIPIcs.

[BG94] Mihir Bellare and Shafi Goldwasser. The complexity of decision versus search. *SIAM Journal on Computing*, 23(1):97–119, 1994.

[BGH$^+$23] Khashayar Barooti, Alex B. Grilo, Loïs Huguenin-Dumittan, Giulio Malavolta, Or Sattath, Quoc-Huy Vu, and Michael Walter. Public-key encryption with quantum keys. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part IV*, volume 14372 of *LNCS*, pages 198–227, Taipei, Taiwan, November 29 – December 2, 2023. Springer, Cham, Switzerland.

[BGS13] Anne Broadbent, Gus Gutoski, and Douglas Stebila. Quantum one-time programs - (extended abstract). In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 344–360, Santa Barbara, CA, USA, August 18–22, 2013. Springer Berlin Heidelberg, Germany.

[BHHP24] John Bostanci, Jonas Haferkamp, Dominik Hangleiter, and Alexander Poremba. Efficient quantum pseudorandomness from Hamiltonian phase states. Cryptology ePrint Archive, Report 2024/1639, 2024.

[BHL$^+$05] Charles H. Bennett, Patrick Hayden, Debbie W. Leung, Peter W. Shor, and Andreas Winter. Remote preparation of quantum states. *IEEE Transactions on Information Theory*, 51(1):56–74, 2005.

[BHMV25] Samuel Bouaziz--Ermann, Minki Hhan, Garazi Muguruza, and Quoc-Huy Vu. On Limits on the Provable Consequences of Quantum Pseudorandomness, 2025. In preparation.

[BJ24] Rishabh Batra and Rahul Jain. Commitments are equivalent to statistically-verifiable one-way state generators. In *65th FOCS*, pages 1178–1192, Chicago, IL, USA, October 27–30, 2024. IEEE Computer Society Press.

[BJSW16] Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. Zero-knowledge proof systems for QMA. In Irit Dinur, editor, *57th FOCS*, pages 31–40, New Brunswick, NJ, USA, October 9–11, 2016. IEEE Computer Society Press.

[BK11] Salman Beigi and Robert König. Simplified instantaneous non-local quantum computation with applications to position-based cryptography. *New Journal of Physics*, 13(9):093036, 2011.

[BK23] James Bartusek and Dakshita Khurana. Cryptography with certified deletion. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 192–223, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland.

[BKS23] James Bartusek, Dakshita Khurana, and Akshayaram Srinivasan. Secure computation with shared EPR pairs (or: How to teleport in zero-knowledge). In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 224–257, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland.

[BMM$^+$25] Amit Behera, Giulio Malavolta, Tomoyuki Morimae, Tamer Mour, and Takashi Yamakawa. A new world in the depths of microcrypt: Separating OWSGs and quantum money from QEFID. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 23–52, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland.

[BMPZ19] Fabio Banfi, Ueli Maurer, Christopher Portmann, and Jiamin Zhu. Composable and finite computational security of quantum message transmission. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part I*, volume 11891 of *LNCS*, pages 282–311, Nuremberg, Germany, December 1–5, 2019. Springer, Cham, Switzerland.

[Bor09] Émile Borel. Les probabilités dénombrables et leurs applications arithmétiques. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 27(1):247–271, 1909.

[BS20a] Nir Bitansky and Omri Shmueli. Post-quantum zero knowledge in constant rounds. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *52nd ACM STOC*, pages 269–279, Chicago, IL, USA, June 22–26, 2020. ACM Press.

[BS20b] Zvika Brakerski and Omri Shmueli. Scalable pseudorandom quantum states. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 417–440, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Cham, Switzerland.

[BW16] Anne Broadbent and Evelyn Wainewright. Efficient simulation for quantum message authentication. In Anderson C. A. Nascimento and Paulo Barreto, editors, *ICITS 16*, volume 10015 of *LNCS*, pages 72–91, Tacoma, WA, USA, August 9–12, 2016. Springer, Cham, Switzerland.

[Can17] Francesco P. Cantelli. Sulla probabilità come limite della frequenza. *Atti Reale Academia Nazionale dei Lincei*, 26(1):39–45, 1917.

[CCS25] Boyang Chen, Andrea Coladangelo, and Or Sattath. The power of a single haar random state: Constructing and separating quantum pseudorandomness. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 108–137, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland.

[CGG24] Kai-Min Chung, Eli Goldin, and Matthew Gray. On central primitives for quantum cryptography with classical communication. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 215–248, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

[CGG$^+$25] Bruno Cavalar, Eli Goldin, Matthew Gray, Peter Hall, Yanyi Liu, and Angelos Pelecanos. On the Computational Hardness of Quantum One-Wayness. *Quantum*, 9:1679, 2025.

[CGGH25] Bruno Pasqualotto Cavalar, Eli Goldin, Matthew Gray, and Peter Hall. A meta-complexity characterization of quantum cryptography. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 82–107, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland.

[CGS02] Claude Crépeau, Daniel Gottesman, and Adam Smith. Secure multi-party quantum computation. In *34th ACM STOC*, pages 643–652, Montréal, Québec, Canada, May 19–21, 2002. ACM Press.

[CK88] Claude Crépeau and Joe Kilian. Achieving oblivious transfer using weakened security assumptions (extended abstract). In *29th FOCS*, pages 42–52, White Plains, NY, USA, October 24–26, 1988. IEEE Computer Society Press.

[CKR11] André Chailloux, Iordanis Kerenidis, and Bill Rosgen. Quantum commitments from complexity assumptions. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *ICALP 2011, Part I*, volume 6755 of *LNCS*, pages 73–85, Zurich, Switzerland, July 4–8, 2011. Springer Berlin Heidelberg, Germany.

[CLM$^+$21] Matthias Christandl, Felix Leditzky, Christian Majenz, Graeme Smith, Florian Speelman, and Michael Walter. Asymptotic Performance of Port-Based Teleportation. *Communications in Mathematical Physics*, 381(1):379–451, 2021.

[CM24] Andrea Coladangelo and Saachi Mutreja. On black-box separations of quantum digital signatures from pseudorandom states. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part III*, volume 15366 of *LNCS*, pages 289–317, Milan, Italy, December 2–6, 2024. Springer, Cham, Switzerland.

[CMS23] Léo Colisson, Garazi Muguruza, and Florian Speelman. Oblivious transfer from zero-knowledge proofs - or how to achieve round-optimal quantum oblivious transfer and zero-knowledge proofs on quantum states. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VIII*, volume 14445 of *LNCS*, pages 3–38, Guangzhou, China, December 4–8, 2023. Springer, Singapore, Singapore.

[DCEL09] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80:012304, 2009.

[DFSS05] Ivan Damgård, Serge Fehr, Louis Salvail, and Christian Schaffner. Cryptography in the bounded quantum-storage model. In *46th FOCS*, pages 449–458, Pittsburgh, PA, USA, October 23–25, 2005. IEEE Computer Society Press.

[DGJ$^+$20] Yfke Dulek, Alex B. Grilo, Stacey Jeffery, Christian Majenz, and Christian Schaffner. Secure multi-party quantum computation with a dishonest majority. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 729–758, Zagreb, Croatia, May 10–14, 2020. Springer, Cham, Switzerland.

[DMS25] Yfke Dulek, Garazi Muguruza, and Florian Speelman. An efficient combination of quantum error correction and authentication. *IACR Communications in Cryptology*, 1(4), 2025.

[DNS10] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Secure two-party quantum evaluation of unitaries against specious adversaries. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 685–706, Santa Barbara, CA, USA, August 15–19, 2010. Springer Berlin Heidelberg, Germany.

[DNS12] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Actively secure two-party evaluation of any quantum operation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 794–811, Santa Barbara, CA, USA, August 19–23, 2012. Springer Berlin Heidelberg, Germany.

[Dru12] Andrew Drucker. New limits to classical and quantum instance compression. In *53rd FOCS*, pages 609–618, New Brunswick, NJ, USA, October 20–23, 2012. IEEE Computer Society Press.

[DS18] Yfke Dulek and Florian Speelman. Quantum ciphertext authentication and key recycling with the trap code. In *13th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC)*, Leibniz International Proceedings in Informatics, pages 1:1–1:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018.

[EGL82] Shimon Even, Oded Goldreich, and Abraham Lempel. A randomized protocol for signing contracts. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *CRYPTO'82*, pages 205–210, Santa Barbara, CA, USA, 1982. Plenum Press, New York, USA.

[EGM24] Francisco Escudero Gutiérrez and Garazi Muguruza. All s_p notions of quantum expansion are equivalent, 2024.

[ELE] ELECTRIC COIN COMPANY. Zcash: Privacy-protecting digital currency.

[FGMR25] Pouria Fallahpour, Alex B. Grilo, Garazi Muguruza, and Mahshid Riahinia. Cryptography from Lossy Reductions: Towards OWFs from ETH, and Beyond. Cryptology ePrint Archive, Report 2025/778, 2025.

[FTH23] Jiani Fei, Sydney Timmerman, and Patrick Hayden. Efficient quantum algorithm for port-based teleportation. *arXiv preprint*, 2023.

[GBO23] Dmitry Grinko, Adam Burchardt, and Maris Ozols. Efficient quantum circuits for port-based teleportation. *arXiv preprint*, 2023.

[GE24] Manuel Goulão and David Elkouss. Pseudo-entanglement is necessary for efi pairs, 2024.

[GH01] Andrew V. Goldberg and Jason D. Hartline. Competitive Auctions for Multiple Digital Goods. In *European Symposium on Algorithms (ESA)*, pages 416–427. Springer Berlin Heidelberg, 2001.

[GJMZ23] Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. In Barna Saha and Rocco A. Servedio, editors, *55th ACM STOC*, pages 1579–1588, Orlando, FL, USA, June 20–23, 2023. ACM Press.

[GLSV21] Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in MiniQCrypt. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 531–561, Zagreb, Croatia, October 17–21, 2021. Springer, Cham, Switzerland.

[GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th ACM STOC*, pages 291–304, Providence, RI, USA, May 6–8, 1985. ACM Press.

[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229, New York City, NY, USA, May 25–27, 1987. ACM Press.

[Gol90] Oded Goldreich. A note on computational indistinguishability. *Information Processing Letters*, 34(6):277–281, 1990.

[Gol06] Oded Goldreich. *Foundations of Cryptography: Volume 1*. Cambridge University Press, USA, 2006.

[Got96] Daniel Gottesman. Class of quantum error-correcting codes saturating the quantum hamming bound. *Physical Review A*, 54:1862–1868, 1996.

[GY25] Alex B. Grilo and Álvaro Yángüez. Quantum pseudoresources imply cryptography, 2025.

[GYZ17] Sumegha Garg, Henry Yuen, and Mark Zhandry. New security notions and feasibility results for authentication of quantum data. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 342–371, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Cham, Switzerland.

[HHH99] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. General teleportation channel, singlet fraction, and quasidistillation. *Physical Review A*, 60:1888–1898, 1999.

[HILL99] Johan HÅstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A Pseudorandom Generator from any One-way Function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

[HLM16] Patrick Hayden, Debbie W. Leung, and Dominic Mayers. The Universal Composable Security of Quantum Message Authentication with Key Recyling, 2016.

[HLM17] Aram W. Harrow, Cedric Yen-Yu Lin, and Ashley Montanaro. Sequential measurements, disturbance and property testing. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1598–1611. SIAM, 2017.

[HM10] Aram Wettroth Harrow and Ashley Montanaro. An efficient test for product states with applications to quantum Merlin-Arthur games. In *51st FOCS*, pages 633–642, Las Vegas, NV, USA, October 23–26, 2010. IEEE Computer Society Press.

[HM24] Taiga Hiroka and Tomoyuki Morimae. Quantum cryptography from meta-complexity. Cryptology ePrint Archive, Report 2024/1539, 2024.

[HMMH$^+$23] Jonas Haferkamp, Felipe Montealegre-Mora, Markus Heinrich, Jens Eisert, David Gross, and Ingo Roth. Efficient unitary designs with a system-size independent number of non-clifford gates. *Communications in Mathematical Physics*, 397(3):995–1041, 2023.

[HMY23a] Minki Hhan, Tomoyuki Morimae, and Takashi Yamakawa. A Note on Output Length of One-Way State Generators and EFIs, 2023.

[HMY23b] Minki Hhan, Tomoyuki Morimae, and Takashi Yamakawa. From the hardness of detecting superpositions to cryptography: Quantum public key encryption and commitments. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 639–667, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland.

[HSS11] Sean Hallgren, Adam Smith, and Fang Song. Classical cryptographic protocols in a quantum world. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 411–428, Santa Barbara, CA, USA, August 14–18, 2011. Springer Berlin Heidelberg, Germany.

[HZB06] Mark Hillery, Mário Ziman, and Vladimír Bužek. Approximate programmable quantum processors. *Physical Review A*, 73(2):022345, 2006.

[IH08] Satoshi Ishizaka and Tohya Hiroshima. Asymptotic teleportation scheme as a universal programmable quantum processor. *Physical review letters*, 101(24):240501, 2008.

[Imp95] Russell Impagliazzo. A personal view of average-case complexity. In *Proceedings of Structure in Complexity Theory. Tenth Annual IEEE Conference*, pages 134–147, 1995.

[Ish15] Satoshi Ishizaka. Some remarks on port-based teleportation. *arXiv preprint*, 2015.

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 126–152, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Cham, Switzerland.

[JLS24] Xiaoyu Ji, Junru Li, and Yifan Song. Linear-communication asynchronous complete secret sharing with optimal resilience. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024,* *Part VIII*, volume 14927 of *LNCS*, pages 418–453, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

[Kil88] Joe Kilian. Founding cryptography on oblivious transfer. In *20th ACM STOC*, pages 20–31, Chicago, IL, USA, May 2–4, 1988. ACM Press.

[KP17] Elham Kashefi and Anna Pappa. Multiparty Delegated Quantum Computing. *Cryptography*, 1(2):12, 2017.

[KPPG19] Aleksander M Kubicki, Carlos Palazuelos, and David Pérez-García. Resource quantification for the no-programing theorem. *Physical review letters*, 122(8):080505, 2019.

[KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In Barna Saha and Rocco A. Servedio, editors, *55th ACM STOC*, pages 1589–1602, Orlando, FL, USA, June 20–23, 2023. ACM Press.

[Kre21] William Kretschmer. Quantum Pseudorandomness and Classical Complexity. In *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, volume 197, pages 2:1–2:20. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.

[KT24] Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 968–978, Vancouver, BC, Canada, June 24–28, 2024. ACM Press.

[Led22] Felix Leditzky. Optimality of the pretty good measurement for port-based teleportation. *Letters in Mathematical Physics*, 112(5):98, 2022.

[Lin13] Yehuda Lindell. A Note on Constant-Round Zero-Knowledge Proofs of Knowledge. *Journal of Cryptology*, 26(4):638–654, 2013.

[LMS22] Alex Lombardi, Fermi Ma, and Nicholas Spooner. Post-quantum zero knowledge, revisited or: How to do quantum rewinding undetectably. In *63rd FOCS*, pages 851–859, Denver, CO, USA, October 31 – November 3, 2022. IEEE Computer Society Press.

[Lo97] Hoi-Kwong Lo. Insecurity of Quantum Secure Computations. *Physical Review A*, 56(2):1154–1162, 1997.

[LT22] Peeter Laud and Riivo Talviste. Review of the state of the art in secure multiparty computation, 2022.

[Lub78] Elihu Lubkin. Entropy of an n-system from its correlation with a k-reservoir. *Journal of Mathematical Physics*, 19(5):1028–1031, 1978.

[LV24] Romi Levy and Thomas Vidick. PRS Length Expansion, 2024.

[LY94] Richard J. Lipton and Neal E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *26th ACM STOC*, pages 734–740, Montréal, Québec, Canada, May 23–25, 1994. ACM Press.

[Maj18] Christian Majenz. Entropy in Quantum Information Theory – Communication and Cryptography. *Ph.D. thesis, Faculty of Science, University of Copenhagen*, 2018.

[Mau11] Ueli Maurer. Constructive Cryptography – A New Paradigm for Security Definitions and Proofs. In *Theory of Security and Applications (TOSCA)*, Lecture Notes in Computer Science, pages 33–56. Springer Berlin Heidelberg, 2011.

[MNY24] Tomoyuki Morimae, Barak Nehoran, and Takashi Yamakawa. Unconditionally secure commitments with quantum auxiliary inputs. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 59–92, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

[MR11] Ueli Maurer and Renato Renner. Abstract Cryptography. In *The Second Symposium on Innovations in Computer Science (ICS 2011)*, pages 1–21. Tsinghua University Press, 2011.

[MS24] Garazi Muguruza and Florian Speelman. Port-Based State Preparation and Applications. *Quantum*, 8:1573, 2024.

[MSSH18] Marek Mozrzymas, Michał Studziński, Sergii Strelchuk, and Michał Horodecki. Optimal port-based teleportation. *New Journal of Physics*, 20(5):053006, 2018.

[MY22] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 269–295, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.

[MY24a] Tomoyuki Morimae and Takashi Yamakawa. One-Wayness in Quantum Cryptography. In Frédéric Magniez and Alex Bredariol Grilo, editors, *19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024)*, volume 310 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 4:1–4:21, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[MY24b] Tomoyuki Morimae and Takashi Yamakawa. Quantum advantage from one-way functions. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part V*, volume 14924 of *LNCS*, pages 359–392, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

[Nao91] Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, January 1991.

[Nay99] Ashwin Nayak. Optimal lower bounds for quantum automata and random access codes. In *40th FOCS*, pages 369–377, New York, NY, USA, October 17–19, 1999. IEEE Computer Society Press.

[NC97] Michael A. Nielsen and Isaac L. Chuang. Programmable Quantum Gate Arrays. *Physical review letters*, 79(2):321–324, 1997.

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum computation and quantum information.* Cambridge University Press, Cambridge, New York, 10th anniversary ed edition, 2010.

[NR06] Moni Naor and Guy N. Rothblum. Learning to impersonate. In *Proceedings of the 23rd International Conference on Machine*

Learning, ICML '06, page 649–656, New York, NY, USA, 2006. Association for Computing Machinery.

[OSP23] Ryan O'Donnell, Rocco A. Servedio, and Pedro Paredes. Explicit orthogonal and unitary designs. In *64th FOCS*, pages 1240–1260, Santa Cruz, CA, USA, November 6–9, 2023. IEEE Computer Society Press.

[PG06] David Pérez-García. Optimality of programmable quantum measurements. *Physical Review A*, 73(5):052315, 2006.

[PG13] Damián Pitalúa-García. Deduction of an upper bound on the success probability of port-based teleportation from the no-cloning theorem and the no-signaling principle. *Physical Review A*, 87(4):040303, 2013.

[PLLP19] Stefano Pirandola, Riccardo Laurenza, Cosmo Lupo, and Jason L. Pereira. Fundamental limits to quantum channel discrimination. *npj Quantum Information*, 5(1), 2019.

[Por17] Christopher Portmann. Quantum authentication with key recycling. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 339–368, Paris, France, April 30 – May 4, 2017. Springer, Cham, Switzerland.

[Pre99] John Preskill. Lecture notes for Physics 219: Quantum computation. *Caltech Lecture Notes*, 1999.

[PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland.

[PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571, Santa Barbara, CA, USA, August 17–21, 2008. Springer Berlin Heidelberg, Germany.

[QDS$^+$19a] Marco Túlio Quintino, Qingxiuxiong Dong, Atsushi Shimbo, Akihito Soeda, and Mio Murao. Probabilistic exact universal quantum circuits for transforming unitary operations. *Physical Review A*, 100(6), 2019.

[QDS$^+$19b] Marco Túlio Quintino, Qingxiuxiong Dong, Atsushi Shimbo, Akihito Soeda, and Mio Murao. Reversing unknown quantum transformations: Universal quantum circuit for inverting general unitary operations. *Physical review letters*, 123(21), 2019.

[Qua20] Willy Quach. UC-secure OT from LWE, revisited. In Clemente Galdi and Vladimir Kolesnikov, editors, *SCN 20*, volume 12238 of *LNCS*, pages 192–211, Amalfi, Italy, September 14–16, 2020. Springer, Cham, Switzerland.

[Rab05] Michael O. Rabin. How to exchange secrets with oblivious transfer. Cryptology ePrint Archive, Report 2005/187, 2005.

[Rom90] John Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd ACM STOC*, pages 387–394, Baltimore, MD, USA, May 14–16, 1990. ACM Press.

[Sat] Or Sattath. Microcrypt Zoo.

[SBZ19] Michal Sedlák, Alessandro Bisio, and Mário Ziman. Optimal probabilistic storage and retrieval of unitary channels. *Physical Review Letters*, 122(17), 2019.

[SKSZ25] Iker Salaberri, Dorota Krajewska, Ekaitz Santazilia, and Eneko Zuloaga, editors. *Investigating Language Isolates: Typological and diachronic perspectives.* John Benjamins Publishing Company, 2025.

[SSMH17] Michał Studziński, Sergii Strelchuk, Marek Mozrzymas, and Michał Horodecki. Port-based teleportation in arbitrary dimension. *Scientific Reports*, 7(1):10871, 2017.

[SV03] Amit Sahai and Salil P. Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM (JACM)*, 50(2):196–249, 2003.

[Ter15] Barbara M. Terhal. Quantum Error Correction for Quantum Memories. *Reviews of Modern Physics*, 87(2):307–346, 2015.

[Unr12] Dominique Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 135–152, Cambridge, UK, April 15–19, 2012. Springer Berlin Heidelberg, Germany.

[Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 755–784, Sofia, Bulgaria, April 26–30, 2015. Springer Berlin Heidelberg, Germany.

[vN28] J. v. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.

[Wat02] John Watrous. imits on the power of quantum statistical zero-knowledge. In *43rd FOCS*, pages 459–470, Vancouver, BC, Canada, November 16–19, 2002. IEEE Computer Society Press.

[Wat09] John Watrous. Zero-Knowledge against Quantum Attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009.

[Wie83] Stephen Wiesner. Conjugate coding. *ACM SIGACT News*, 15(1):78–88, 1983.

[Wil13] Mark M Wilde. *Quantum information theory.* Cambridge university press, 2013.

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

[WW06] Stefan Wolf and Jürg Wullschleger. Oblivious transfer is symmetric. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 222–232, St. Petersburg, Russia, May 28 – June 1, 2006. Springer Berlin Heidelberg, Germany.

[Yan22] Jun Yan. General properties of quantum bit commitments (extended abstract). In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 628–657, Taipei, Taiwan, December 5–9, 2022. Springer, Cham, Switzerland.

[Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd FOCS*, pages 160–164, Chicago, Illinois, November 3–5, 1982. IEEE Computer Society Press.

[YAVV22] Vijay Kumar Yadav, Nitish Andola, Shekhar Verma, and S. Venkatesan. A Survey of Oblivious Transfer Protocol. *ACM Computing Surveys*, 54:211:1–211:37, 2022.

[YRC20] Yuxiang Yang, Renato Renner, and Giulio Chiribella. Optimal Universal Programming of Unitary Gates. *Physical review letters*, 125(21):210501, 2020.

Abstract

In this thesis we explore multiple angles of communication between parties that have access to quantum resources, such as channels, qudits, and computers.

We start by studying the options to set up a communication link, in other words, we explore the relation between different resources needed for establishing a channel capable of transmitting quantum information. In particular, we consider a scenario where the sender knows a classical description of the qudit they intend to send, and the receiver's operations are restricted to classical ones. Our main result is that the accuracy of the transmission scales inverse exponentially with the number of pre-shared entangled qudits that the sender and receiver share.

We later ask what extra properties we could expect from general quantum channels. In this direction, we give a protocol for authenticating a noisy channel, that is, a set of rules that two communicating parties can follow in order to detect if the changes that happening to qudits when sent through a quantum channel are a consequence of naturally occurring noise or undesired tampering by a third party. Moreover, we prove that our protocol requires access to poly-logarithmic fewer qubits than the previously known techniques.

For the rest of the dissertation we assume that the parties have access to ideal quantum communication channels and look at what these channels could be useful for. Our first result is a round-optimal quantum protocol for oblivious transfer: the goal of this task is for two parties to share one message out of two, with one of the parties choosing which one (and only one) and the other providing the chosen message, without the latter party knowing which message out of the two was chosen by the former. This type of abstraction of tasks to their most basic form allows them to be used as building blocks for

more complex protocols, as long as the results are proven under a composable notion of security, and the set-up assumptions are clear. In particular, although our oblivious transfer protocol allows for instantiations both in the plain and quantum random oracle models (by basically lifting the properties of an underlying zero-knowledge protocol), we obtain round optimality in the quantum random oracle model.

Some set-up assumptions are necessary to be able to prove the existence of almost any cryptographic primitive. One fascinating property of communication under the laws of quantum mechanics is the possibility of security under assumptions weaker than classical pseudorandomness, that is, the existence of algorithms that extend the length of random bit-strings to random looking (for computationally bounded observers) bit-strings, whose existence would in particular imply that $\mathsf{P} \neq \mathsf{NP}$. The last 3 chapters of this dissertation are an exploration of quantum pseudorandomness; in this case we have random (looking) quantum states instead of bit-strings. We first show that quantum pseudorandomness behaves very differently from classical pseudorandomness because there is an inherent impossibility of shrinking the size of a pseudorandom object. Finally, we give conditions on promise problems for the existence of quantum pseudorandom primitives. Namely, that they admit a quantum reduction that loses information about its input.

Samenvatting

In dit proefschrift onderzoeken we verschillende communicatiemogelijkheden tussen partijen die toegang hebben tot kwantumbronnen; zoals kanalen, qudits en computers.

We beginnen met het bestuderen van de mogelijkheden om een communicatieverbinding op te zetten; met andere woorden, we onderzoeken de relatie tussen de verschillende bronnen die nodig zijn om een kanaal tot stand te brengen dat kwantuminformatie kan verzenden. We beschouwen met name een scenario waarin de zender een klassieke beschrijving kent van de qudit die zij wil verzenden, en de ontvanger zich tot klassieke bewerkingen beperkt. Onze belangrijkste uitkomst is dat de nauwkeurigheid van de transmissie exponentieel omgekeerd evenredig is met het aantal vooraf gedeelde verstrengelde qudits dat de zender en de ontvanger delen.

Vervolgens vragen we ons af welke extra eigenschappen we van algemene kwantumkanalen kunnen verwachten. In deze richting presenteren we een protocol voor het authenticeren van een ruisig kanaal, dat wil zeggen een reeks regels die twee communicerende partijen kunnen volgen om te detecteren of de veranderingen die optreden in qudits wanneer ze via een kwantumkanaal worden verzonden, het gevolg zijn van natuurlijk voorkomende ruis of ongewenste beïnvloeding door een derde partij. Bovendien bewijzen we dat ons protocol toegang vereist tot poly-logaritmisch minder qubits dan de eerder bekende technieken.

Voor de rest van het proefschrift gaan we ervan uit dan de partijen toegang hebben tot ideale kwantumcommunicatiekanalen en bekijken we waarvoor deze kanalen nuttig zouden kunnen zijn. Ons eerste resultaat is een rondoptimaal kwantumprotocol voor *oblivious transfer*: het doel van deze taak is dat twee partijen één bericht uit twee delen, waarbij één van de partijen kiest welke (en slechts één) en de andere het gekozen bericht levert, zonder dat de laatste partij weet welk bericht uit twee door de eerste is gekozen. Dit type abstractie van taken tot hun meest basale vorm maakt het mogelijk dat ze worden gebruikt als bouwstenen voor complexere protocollen, zolang de resultaten bewezen worden onder een samenstelbaar begrip van beveiliging en de opstellingen duidelijk zijn. Hoewel ons oblivious-transfer protocol instantiaties toestaat in zowel het gewone als het kwantum-random-orakelmodel (door in feite de eigenschappen van een onderliggend zero-knowledge-protocol op te heffen), verkrijgen we rondoptimaliteit in het kwantum-random-orakelmodel.

Sommige aannames zijn nodig om het bestaan van vrijwel elk cryptografisch primitief te kunnen bewijzen. Een fascinerende eigenschap van communicatie volgens de wetten van de kwantummechanica is de mogelijkheid van beveiliging onder aanames die zwakker zijn dan klassieke pseudorandomness, dat wil zeggen het bestaan van algoritmen die de lengte van willekeurige bit-strings uitbreiden naar willekeurig uitziende (voor computationeel begrensde waarnemers) bit-strings, waarvan het bestaan met name zou impliceren dat $\mathsf{P} \neq \mathsf{NP}$. De laatste drie hoofdstukken van dit proefschrift zijn een verkenning van kwantumpseudorandomness; in dit geval hebben we willekeurige (uitziende) kwantumtoestanden in plaats van bit-strings. We laten eerst zien dat kwantumpseudorandomness zich heel anders gedraagt dan klassieke pseudorandomness, omdat het inherent onmogelijk is om de grootte van een pseudorandom object te verkleinen. Ten slotte geven we voorwaarden voor promise-problemen voor het bestaan van kwantumpseudorandom primitieven. Namelijk dat ze een kwantumreductie toelaten die informatie over de invoer verliest.

Acknowledgements

I would like to start by thanking Florian, my supervisor, for being such a wonderful colleague. Thank you for your patience, care and support. For believing that there are many ways of doing science and helping me find mine.

I thank my advisors Chris and Paola for the guidance and the structure you provided. I really appreciate your willingness to make the university safer for minoritized communities. I would also like to thank my committee for evaluating my thesis.

I am immensely grateful to my collaborators, after the loneliness of the pen and paper I knew I could count on your cheerfulness — regardless of our time zones. Florian, Huy, Samuel, Minki, Gina, Alex, Léo, thank you. Pouria and Mahshid, the two weeks we shared in Lyon are the most fun I ever had doing research, thank you for welcoming me.

The commute all the way to Amsterdam for four years was worth the days that I got to see the friends I made there: Jonas, Freek, Jana, Koen, Clara, Said, Salvatore, Lorenzo, Fran, Jelena, Subha... you have given meaning to my PhD, thank you for caring so much. I want to also thank the people at the margins of our university for not closing your eyes in face on genocide, for teaching me that colonial history is very present in our research, and for all the efforts put into imagining better futures together.

But if I chose to remain living in Delft it is because, as my parents claim, it has become my home. This home abroad has been built with the help of many people,

the running group and foodsharing, thank you for the freedom,
Gustavo, thank you for the fire (and the shadows),
Rishabh, thank you for the stories,

Garoa, Amaia and Mireia, eskerrikasko erroak saretuz luzatzeagatik,
Alexandra, thank you for the knowledge you cannot find in books,
Pablo and Ralf, gracias por la espontaneidad,
Kevin and Kruti, thank you for every shared meal,
Uri, gràcies per l'amor i el temps,
Ankita, my soulmate, thank you for helping me understand the world,
Aniket, mazha moramba, majhya sobat rahayla saathi dhanyabad,
and the ones who were once here and left, but we managed to stay in each other's life: Jin, Reuben, Malu and Sparsh, thank you for warming my heart with your messages.

Finally, I am so lucky to have Irati and Jelena both as friends and as paranymphs. Irati, bizitzako momentu garrantzitsu guztitan nere ondon nahi zaitudalako, eskerrikasko nere ondon hazteatik urte haue guztitan, hasieratik. Jelena, because I walk around searching for postcards for (y)our office, thank you for making life so interesting.

I want to dedicate this thesis to my family, because without their unconditional support I would not have been able to complete it. Amatxo, aitatxo, eskerrikasko babes eta askatasun guztiyeatik, asko maite zaituztet.

Delft, September 2025