Master's thesis in Mathematics

Oblivious Transfer in Quantum Cryptography

Simon S. Erfurth
Advisor: Joan Boyar

1st June, 2021

DEPARTMENT OF MATHEMATICS
AND COMPUTER SCIENCE

PRINT: GRAFISK CENTER, SDU

Resumé

Kvantecomputere bliver stadig mere anvendelige i praksis, og kvantekommunikation bliver allerede anvendt i praksis. Dermed stiger også relevansen af dettes speciales fokus; kvantekryptografi. Vi introducerer kvantemekanik, kvantekomputabilitetsteori, og kvanteinformationsteori og analyserer prøveudtagelsesspil på kvanterum. Ved hjælp af disse værktøjer viser vi, at BBCS92 protokollen er en sikker implementation af den kryptografiske primitiv "oblivious transfer" overfor vilkårlige kvantemodstandere. I sig selv bruger denne protokol ikke andre kvanteredskaber end kvantekommunikation. Endeligt så viser vi, hvordan man kan konstruere en protokol for "bit-commitments" fra kvantesikre en-vejs funktioner, hvilket medfører, at BBCS92 protokollen er i MiniQCrypt, og dermed er der også protokoller for sikker distribueret beregning af arbitrære funktioner i MiniQCrypt.

Abstract

With quantum computing becoming ever more usable in practise, and quantum communication already being feasible for practical applications, quantum cryptography is also becoming increasingly relevant. We introduce quantum mechanics, quantum computation theory, and quantum information theory, and investigate sampling games in a quantum setting. With this toolkit we show that the BBCS92 protocol for the cryptographic primitive oblivious transfer is secure against arbitrary quantum adversaries. On its own this protocol uses no quantum resources, other than quantum communication. Finally, we show how to create bit-commitments from post-quantum one-way functions, implying that the BBCS92 protocol is in MiniQCrypt, and hence, so are protocols for arbitrary secure multi-party computations.

Contents

1 Introduction

Over the last 100 years, many discoveries and widely believed theories in physics have made it increasingly clear that the world we live in is governed by quantum forces, rather than "classical" forces, as has been assumed for thousands of years. The postulates of quantum mechanics describes how the theories describing this "new" quantum world should be formed, and are widely accepted. It should come as no surprise that the postulates of quantum mechanics also influences what can and cannot be done by a computer, and indeed, since the 1980s quantum computing has been studied. That there is a difference between what can be achieved on a quantum computer and a classical computer became extremely clear in 1994, when Peter Shor developed a quantum algorithm for factoring an integer in polynomial time, a problem that it is generally believed to be impossible to solve efficiently on a classical computer [NC02].

The discovery of this algorithm — referred to as Shor's Algorithm — indicates the impact that quantum computers will have on cryptography. With the ability to factor integers in polynomial time, many common and widely used cryptographic constructions are no longer secure, eg. RSA. Hence, it is of vital importance to study how quantum computers, and quantum mechanics in general, affect cryptography [NC02]. This is the focus of *quantum cryptography*, the field that this thesis falls under. Quantum cryptography can generally be separated into two categories; post-quantum cryptography, where one investigates which (classical) problems appears to still be hard for quantum computers, and a more general category where one also uses quantum resources, allowing one to obtain results that are not possible in classical cryptography. This category is also referred to as quantum cryptography. The crown jewel of this last category Bennett and Brassard's discovery of a protocol that, by using quantum communication, achieves unconditional secure key-exchange, which is a problem where a classical solution has to rely on a computational assumption [BB84]. Quantum communication is the ability to prepares, exchange, and measure quantum states. This is much easier to do than general quantum computation, and is feasible in practise [Dia+16].

In this thesis, we focus on the cryptographic primitive oblivious transfer. There are many varieties of oblivious transfer (Rabin-OT, $GF(q)$–OLFE, and $\binom{n}{k}$-OT), but all of these can be implemented using $\binom{2}{1}$-OT as a primitive [Cré87; IK97; WW06], so we consider only this variant. $\binom{2}{1}$-OT (from now on just referred to as oblivious transfer or OT) consists of two parties, one of which has two messages that the other party should learn only one of, but which one they learn is up to them. The other party should remain ignorant of which message was chosen.

Oblivious transfer is interesting and relevant to study, due to its widespread use in protocols for secure multi-party computation (MPC). A secure multi-party computation consists of two or more parties, who despite not trusting

each other, still want to compute a function based on their joint input. It is well known that secure implementations of oblivious transfer and bit-commitment, implies the existence of secure protocols for evaluating arbitrary (classical) functions [CDN15]. More recently it has been shown that if the implementations are secure in the quantum setting, so are the protocols for MPC [IPS08; Unr10], which in turn imples the existence of secure protocols for evaluating two-party quantum circuits [DNS12].

Using only classical tools, it is widely believed that one needs a computational assumption to implement oblivious transfer, much like one needs for key-exchange. It is therefore only natural to ask if it is possible to implement oblivious transfer with unconditional security using quantum resources? Unfortunately, the answer turns out to be *no* [Lo97]. It is, however, possible to securely implement oblivious transfer using quantum communication and post-quantum one-way functions, which are functions that are easy to evaluate, but hard to find pre-images of, even with quantum resources. These are generally believed to exists, with numerous candidates available [Gri+20]. Taking inspiration from Impagliazzo's five worlds [Imp95], we call a world where post-quantum one-way functions exists — in addition to quantum communication and computation — for *MiniQCrypt*.

Towards showing that oblivious transfer is in MiniQCrypt, we start out by introduce quantum mechanics, the framework for working with quantum resources, in Section 3.1. We generally limit our considerations to qubits, which is the quantum analogy of a classical bit and hence the focus of quantum computing and cryptography. In Section 3.2 we give a very short overview of quantum computation theory, and in Section 3.3 of quantum information theory. Both sections focus on a select few topics that are of special relevance for quantum cryptography. In Section 2.3 and 3.4 we introduce and analyse sampling games in first a classical and then a quantum setting. In section 4 we start our investigation of quantum cryptography, with some fundamental examples, such as the earlier mentioned protocol for quantum key-distribution, and the notion of security that we use. This leads us to Section 5 where we turn our attention to oblivious transfer, first describing the protocol for oblivious transfer in Section 5.1 and then proving that it is secure assuming the existence of ideal bit-commitments in Section 5.2. To show this we use the theory for sampling strategies that we developed in Section 3.4. Finally, in section 5.3 we see how one can construct bit-commitments from post-quantum one-way functions, allowing us to conclude that oblivious transfer and MPC is in MiniQCrypt.

2 Preliminaries and classical results

In this section we present some mathematical and cryptographic preliminaries, that we will use later in the thesis. In Section 2.3 we also introduce and develop some results for sampling strategies in a classical setting.

2.1 Mathematical preliminaries

2.1.1 Notation

We will use the following (slightly) non-standard notation.

- For a natural number $n \in \mathbb{N}$ we use $[n]$ to denote the set of natural numbers up to and including n, i.e. $[n] = \{1, 2, \ldots, n-1, n\}$.
- For specifying a subsequence of a sequence $a = (a_1, a_2, \ldots, a_n)$ we may write a_J with $J \subseteq [n]$. That is, if $J = \{i_1, i_2, \ldots, i_j\} \subseteq [n]$ then $a_J = (a_{i_1}, a_{i_2}, \ldots, a_{i_j})$.
- For any finite set S we use the notation $s \leftarrow_R S$ to denote that s was chosen uniformly at random from the elements in S.
- To avoid changing notation, we use bra-ket notation throughout. This notation is typically used in quantum mechanics. Summarising this notation; $|v\rangle$ is a vector, $\langle w|$ is the functional mapping $|v\rangle \mapsto \langle w|v\rangle$ where $\langle v|w\rangle$ denotes the inner product between $|v\rangle$ and $|w\rangle$. $|v\rangle\langle w|$ denotes the outer product of $|v\rangle$ and $|w\rangle$.

  If M is a linear map then we note that $(\langle v|M)|w\rangle = \langle v|(M|w\rangle)$, and will therefore just write $\langle v|M|w\rangle$. Note that $\langle v|M$ is the composition of the two linear maps, and is itself a linear functional.

  In the case where we are working in $\mathbb{C}^n$ with the Euclidean inner product, one can think of $\langle v|$ as $|v\rangle^T$, in which case all of this notation is just straight forward matrix multiplication.

2.1.2 Linear algebra

While we assume some familiarity with elementary linear algebra, we present some slightly more advanced definitions and results in this section. Throughout this section we let $(V, \langle\cdot|\cdot\rangle)$ be a finite dimensional Hilbert space,[1] and denote by $\mathbb{B}(V)$ the set of linear maps (or operators) on V, which is itself a Hilbert space. We denote the identity operator on V by $\mathbb{I}_V$. When V is clear from the context, we will omit the subscript. Alternatively, for any vector $|v\rangle$ we may write $\mathbb{I}_{|v\rangle}$ to emphasise that $\mathbb{I}$ is the identity on the space that $|v\rangle$ is in.

[1] We note that oftentimes we only need V to be a vector space, but to avoid using additional non-essential terms, we will just always assume that V is a Hilbert space, and refer to it as such.

Adjoint For a map $T \in \mathbb{B}(V)$ we define the adjoint of T, denoted T^*, to be the unique operator on V satisfying that all $|v\rangle, |w\rangle \in V$ we have that

$$\langle v|T^*|w\rangle = \overline{\langle v|T|w\rangle}.$$

When working with $\mathbb{C}^n$ we note that $\mathbb{B}(\mathbb{C}^n) = \mathbb{M}_n(\mathbb{C})$, and T^* is just the conjugated transpose of T.

Unitary An operator $U \in \mathbb{B}(V)$ is said to be unitary if $U^*U = UU^* = \mathbb{I}$.

Projections A projection on V is a map $T \in \mathbb{B}(V)$ such that for all vectors $|v\rangle \in V$ we have that $T^2|v\rangle = T|v\rangle$. We note that for any normal vector $|w\rangle \in V$ the map $P_w \in \mathbb{B}(V)$ defined by $|v\rangle \mapsto \langle v|w\rangle|w\rangle$ is a projection.

For any orthonormal basis $B = \{|b_1\rangle, \ldots, |b_n\rangle\}$ for V we can decompose any vector $|v\rangle \in V$ as

$$|v\rangle = \sum_{|b\rangle \in B} P_b|v\rangle = \sum_{|b\rangle \in B} \langle v|b\rangle|b\rangle.$$

Additionally, if $W \subseteq V$ is a subspace, we define the projection of $|v\rangle$ onto W by

$$P_W|v\rangle := \sum_{|b\rangle \in B_W} P_b|v\rangle,$$

where B_W is any orthonormal basis for W. Hence, we can also decompose $|v\rangle$ in terms of W and its orthogonal compliment;

$$|v\rangle = P_W|v\rangle + P_{W^\perp}|v\rangle.$$

A set of projections, $\{P_m\}_{m \in M}$, is said to be a set of orthogonal projections if each P_m is a projection and $P_m P_n = 0$ whenever $m \neq n$.

Tensor products Tensor products are a convenient way to combine Hilbert spaces to form some larger Hilbert space. Assume that $(V, \langle\cdot|\cdot\rangle_V)$ and $(W, \langle\cdot|\cdot\rangle_W)$ are Hilbert spaces with dimension n and m, respectively, over a common field $\mathbb{F}$. Then $V \otimes W$ is a Hilbert space of dimension nm over $\mathbb{F}$, and $V \otimes W$ consists of linear combinations of tensor products of elements from V and W. More concretely, if $\{|v_i\rangle\}_{i \in i}$ is a basis for V and $\{|w_j\rangle\}_{j \in J}$ is a basis for W, then $\{|v_i\rangle \otimes |w_j\rangle\}_{i \in I, j \in J}$ is a basis for $V \otimes W$, and $V \otimes W$ consists of elements of the form

$$\sum_{i,j} \alpha_{ij}|v_i\rangle \otimes |w_j\rangle.$$

We define an inner product on $V \otimes W$ by defining it on the basis elements on $V \otimes W$ by

$$\langle(|v\rangle \otimes |w\rangle)|(|v'\rangle \otimes |w'\rangle)\rangle = \langle v|v'\rangle_V \cdot \langle w|w'\rangle_W,$$

and extending it by linearity to the entirety of $V \otimes W$. Doing this, we obtain a well-defined inner product on $V \otimes W$.

Trace For a square matrix $A = (a_{ij})$ we define the trace of A to be the sum of the entries on the diagonal;

$$\operatorname{tr} A := \sum_i a_{ii}.$$

The trace is linear, but more importantly, it is invariant under cyclic permutations, i.e. if A is a $m \times n$ matrix, B is a $n \times k$ matrix, and C is a $k \times m$ matrix, then

$$\operatorname{tr}(ABC) = \operatorname{tr}(BCA) = \operatorname{tr}(CAB).$$

Importantly, this also show that the trace is invariant under change of basis, since $\operatorname{tr}(P^{-1}AP) = \operatorname{tr}(PP^{-1}A) = \operatorname{tr} A$. Hence, we can define the trace of an arbitrary linear operator on finite dimensional Hilbert spaces by identifying the operator with any matrix representation of it.

We emphasise the following special case of the cyclic invariance. Let $|v\rangle, |w\rangle \in V$. Then the tensor product of the outer product $|v\rangle\langle w|$ is equal to the inner product $\langle w|v\rangle$, since

$$\operatorname{tr}(|v\rangle\langle w|) = \operatorname{tr}\langle w|v\rangle = \langle w|v\rangle.$$

In the case where A is the tensor product of two operators, $A = T \otimes S$, the trace of A is given by

$$\operatorname{tr} A = \operatorname{tr} T \cdot \operatorname{tr} S.$$

Additionally, we define the *partial trace* of A with respect to T (equivalently S) to be

$$\operatorname{tr}_T A = \operatorname{tr}(T) \cdot S,$$

and we note that this can be extended by linearity to sums of tensors. If T is in a system named $\mathcal{T}$, we may also use $\operatorname{tr}_{\mathcal{T}} A$ to mean the same as $\operatorname{tr}_T(A)$.

Positive semi-definite An operator $T \in \mathbb{B}(V)$ is said to be positive semi-definite, or just positive, if for all $|v\rangle \in V$

$$\langle v|T|v\rangle \geq 0.$$

We write this as $T \geq 0$.

Cauchy-Schwarz and Titu's Lemma The Cauchy-Schwarz inequality states that for all vectors $|u\rangle, |v\rangle$ of a Hilbert space it is the case that

$$|\langle u|v\rangle|^2 \leq \langle u|u\rangle\langle v|v\rangle.$$

A direct corollary is Titu's Lemma.

Corollary 2.1 *Titu's Lemma.*
For $\{u_i\}_{i\in[n]}, \{v_i\}_{i\in[n]}$ sequences of positive real numbers the following inequality holds.

$$\frac{\left(\sum_{i=1}^n u_i\right)^2}{\sum_{i=1}^n v_i} \leq \sum_{i=1}^n \frac{u_i^2}{v_i}.$$

Proof. Let $|v\rangle$ and $|u\rangle$ be vectors in $\mathbb{R}^n$ with the following entries,

$$|u\rangle := \left(\frac{u_1}{\sqrt{v_1}}, \frac{u_2}{\sqrt{v_2}}, \ldots, \frac{u_n}{\sqrt{v_n}}\right),$$
$$|v\rangle := (\sqrt{v_1}, \sqrt{v_2}, \ldots, \sqrt{v_n}).$$

Then the Cauchy-Schwarz inequality implies that

$$\begin{aligned} |\langle u|v\rangle|^2 &\leq \langle u|u\rangle\langle v|v\rangle \\ \iff \left(\sum_{i=1}^n u_i\right)^2 &\leq \left(\sum_{i=1}^n \frac{u_i^2}{v_i}\right)\left(\sum_{i=1}^n v_i\right) \\ \iff \frac{\left(\sum_{i=1}^n u_i\right)^2}{\sum_{i=1}^n v_i} &\leq \sum_{i=1}^n \frac{u_i^2}{v_i}, \end{aligned}$$

which was what was wanted. □

2.1.3 Probability theory

Here we introduce standard definitions and notation from probability theory with some formality, but for most of the project we omit these formalities. Let $(\Omega, \Pr)$ be a probability space. Then a random variable is a function $X\colon \Omega \to \mathcal{X}$ where $\mathcal{X}$ is an arbitrary finite set. In this project we will usually have $\mathcal{X} = \{0,1\}$. The distribution of X, P_X, is given by

$$P_X(x) = \Pr[X = x] = \Pr[\{\omega \in \Omega\colon X(\omega) = x\}],$$

where we have used x to indicate a specific outcome of the random variable X. This will be standard notation whenever we are dealing with random variables. The joint distribution of two (or more) random variables, P_{XY}, is then

$$P_{XY}(x,y) = \Pr[X = x \wedge Y = y],$$

and the random variables X and Y are said to be independent if $P_{XY}(x,y) = P_X(x)P_Y(y)$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. The expected value of a random variable is defined as

$$\mathrm{E}[X] = \sum_{x\in\mathcal{X}} x \cdot P_X(x).$$

Additionally, we will make use of the concave version of Jensen's inequality, which a proof of can be found in [Sch17].

Theorem 2.2 *Jensen's inequality.*
If X is a random variable and f is concave function, then

$$\mathrm{E}[f(X)] \leq f(E[X]).$$

Finally, we will use Shannon's notion of the entropy for how much randomness there is left in a random variable

Definition 2.3 *Shannon entropy.*
If X is a random variable, the entropy about the outcome of X is

$$\mathrm{H}(X) := -\sum_{x \in \mathcal{X}} P_X(x) \lg P_X(x).$$

In the special case where $\mathcal{X} = \{0,1\}$, and X takes on the value 0 with probability p, we define the binary entropy function h as

$$\mathrm{h}(x) = -p \lg p - (1-p) \lg (1-p).$$

2.1.4 Strings

In this section we present some definitions and results about strings. For notation, we will refer to the i'th entry in a string b as b_i, i.e. we think of b as $b = (b_1, b_2, \ldots, b_n) \in \{0,1\}^n$ for some $n \geq 0$.

Definition 2.4 *(Relative) Hamming weight.*
For a string $b \in \{0,1\}^n$, the Hamming weight of b is the number of non-zero entries in b. Formally we define the Hamming weight of b to be

$$\mathrm{wt}(b) := |\{i \in [n] \colon b_i = 1\}|,$$

and the relative Hamming weight of b as

$$\omega(b) = \frac{\mathrm{wt}(b)}{n}.$$

By convention, the relative Hamming wight of the empty string is defined to be 0.

For a string b with relative Hamming weight $\omega(b)$ we denote the random variable obtained by picking an entry from b, uniformly at random, by B, and note that the expected value of B is $\mathrm{E}(B) = \omega(b)$.

For estimating the relative Hamming weight of a string, Hoeffding's inequality will prove to be very useful. In general, Hoeffding's inequalities are a classic result from probability theory, which provides an upper bound for the probability that a sum of bounded random variables from the same distribution deviates from their expected value by more than a certain amount. For a proof we refer to Hoeffing's original paper from 1963 [Hoe63]. We state it here for

random samples from a bit-string without replacement, and in terms of its relative Hamming weight.

Theorem 2.5 *Hoeffding's inequality.*
Let $b \in \{0,1\}^n$ be a bit-string with relative Hamming weight $\mu = \omega(b)$. Let the random variables $X_1, X_2, \ldots, X_k$ be obtained by sampling $k \leq n$ random entries from b without replacement. Then for any $\delta > 0$ the random Variable $\overline{X} := \frac{1}{k}\sum X_i$ satisfy

$$\Pr\left[\left|\overline{X} - \mu\right| \geq \delta\right] \leq 2\exp(-2\delta^2 k).$$

2.2 Cryptographic preliminaries

In this section we will first introduce a few results from classical cryptography. Then, in Section 2.2.1, we spend some time introducing the notion of security that is typically used for classical multi-party computation.

The first result that we introduce is the classical result that a one-time-pad obtains perfect security. We will implicitly use this result when arguing that if some value is (close to) uniformly random, then XOR'ing a message with it gives a value that is itself (close to) uniformly random.

Theorem 2.6
Let $M \in \{0,1\}^n$ be a message of length n, and suppose that a key K is drawn uniformly at random from $\{0,1\}^n$. For an adversary that knows scheme and spaces, but does not know the value of K, the ciphertext $C = M \oplus K$ reveal no information about M.

We will make use of one-way functions, which are functions that are easy to evaluate on a given input, but hard to find the pre-image of. Note that the existence of one-way functions has not been proven, but is a common assumption, with numerous candidates available. Formally we define one-way functions as follows.

Definition 2.7 *One-way functions (OWF).*
A function $f : \{0,1\}^* \to \{0,1\}^*$ is said to be one-way if f can be evaluated in polynomial time, but all polynomial-time adversaries $\mathcal{A}$ have only negligible chance of finding a pre-image of f. That is, for all sufficiently large n, and for all polynomial-time $\mathcal{A}$

$$\Pr[f(\mathcal{A}(f(x))) = f(x)]$$

is negligible in n, where the probability is taken over a uniformly random choice of $x \in \{0,1\}^n$.

Technically, the next result we present makes use of definitions from Section 2.2.1 and Section 2.2.2. However, we state it here using only layman terms, and note that the technical rewording is mainly the use of the hybrid-model, and is completely straight forward.

Theorem 2.8
Assuming the existence of a protocol that securely implements oblivious transfer, there exists protocols that securely implements every classical multi-party functionality F

2.2.1 Relevant security definitions

In this section we introduce the relevant notion of security in the classical setting. We will later extend this notion to the quantum setting. The work here is based on earlier projects, but the canonical reference is Goldreich's Foundations of Cryptography [Gol04].

Before we start, we stress that this is the canonical definition of security for secure multipart computation (adapted to a two-party setting), and hence it is often meaningless for other tasks, such as key exchange, or (a)symmetric encryption. There exists meaningful notions of security for these problems, but since they are not the focus of this project, we will not include them in any formal way, but introduce the intuition behind, as needed.

The family of problems that we consider is creating protocols for functionalities. A functionality, $\mathcal{F}$, receives input from multiple parties, and evaluates a function on their joint input, resulting in a distinct output for each party. However, the parties does not trust each other, so they cannot just share their input with each other, nor is there some third party that they all trust. A protocol should describe how the parties can evaluate the functionality such that they obtain their output, without a trusted party. A protocol is said to enjoy correctness if the output of the parties evaluating the protocol as described, is the same as the output they would obtain by using a theoretical trusted party. Additionally, a protocol is secure if it has the property that no party learn anything they are not supposed to learn, regardless of how they act. We are only interested in secure protocols, that enjoy correctness.

We use the phrase "not supposed to learn", since even with a trusted party, they will of course learn something, eg. that the other party showed up and participated. Essentially, we will argue that a two-party protocol $\Pi = (A, B)$ securely implements $\mathcal{F}$ in the *real-world*, by comparing the real-world execution of Π to an execution in the *ideal-world*, where we assume that there *is* a trusted party that will take care of running $\mathcal{F}$, and only give the parties their respective output of $\mathcal{F}$. Abusing notation, we will also refer to this trusted part as $\mathcal{F}$. Security in the ideal-world follows from construction; we assume the trusted party can be trusted, and only give the parties the output they are supposed to obtain from $\mathcal{F}$. Hence it will only leak what must *inherently* be leaked by $\mathcal{F}$.

Towards formally defining this notion of security, we specify that the functionality $\mathcal{F}$ is a classical (polynomial time) interactive machine, that specifies the instructions to realise a task. The (efficient) two-party protocol $\Pi = (A, B)$ consists of two classical (polynomial time) interactive machines, A and B. An adversary $\mathcal{A}$ is an arbitrary classical (polynomial time) interactive machine,

intending to attack a protocol. Henceforth, we will only specify when something is not in polynomial time, since we are generally interested in efficient protocols.

When considering a real-world execution in the presence of an adversary we will assume that the adversary $\mathcal{A}$ statically corrupts one of the parties A or B (for the rest of this section we assume A for notation, but the definitions are symmetric for B), end therefore have full control over the corrupted party, whom we might therefore refer to as either the corrupted party or, abusing notation, just $\mathcal{A}$. This means that $\mathcal{A}$ sees everything that the corrupted party sees, including its input, and that $\mathcal{A}$ can specify every action that the corrupted party takes, including the messages it sends, and what it should output at the end of the protocol.

We can now formally define computational security against real-world adversaries by the existence of a polynomial time simulator Sim in the ideal-world, such that the joint output of the honest party and the adversary in the real-world, $\text{Out}_{\Pi,\mathcal{A},B}$, is computationally indistinguishable from the joint output of the honest party and the simulator in the ideal-world, $\text{Out}_{\mathcal{F},\text{Sim},\hat{B}}$, where $\hat{B}$ is used to indicate the "dummy" version of B, that simply passes its input to $\mathcal{F}$, and output whatever output $\mathcal{F}$ sends it.

Definition 2.9 *Computational security against corrupted A.*
A protocol $\Pi = (A, B)$ implements a classical functionality $\mathcal{F}$ with computational security against corrupted A, if for any polynomial-time real-world adversary $\mathcal{A}$ there exists an efficient ideal-world simulator Sim such that for any input to A and B it holds that the joint output in the real- and ideal-world are computationally indistinguishable, i.e.

$$\text{Out}_{\Pi,\mathcal{A},B} \overset{c}{\approx} \text{Out}_{\mathcal{F},\text{Sim},\hat{B}}.$$

Equivalently, we say that there is no distinguisher D that can distinguish between the situation pictured in Figure 1 and the situation pictured in Figure 2, with non-negligible advantage.

If a protocol is secure against both corrupted A's and corrupted B's, we will just say that it is secure.

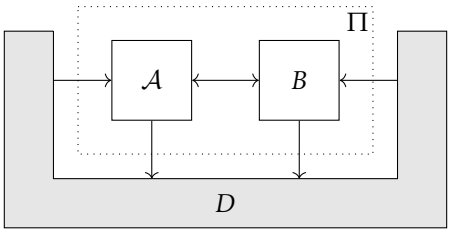

Figure 1: Illustration of the interactions in the real-world.

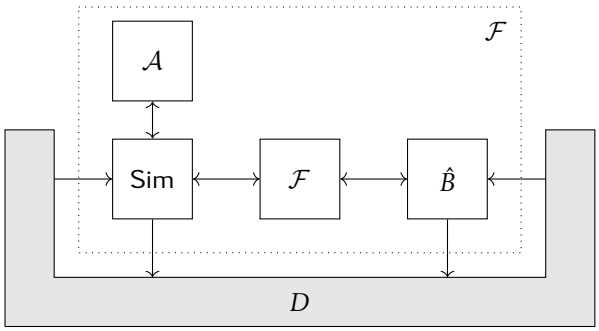

Figure 2: Illustration of the interactions in the ideal-world. Note that we have included an interaction between Sim and $\mathcal{A}$, to illustrate that Sim is allowed to use $\mathcal{A}$ as a subroutine.

Remark 2.10
We have omitted considering an important, but easily avoidable, problem with our definition. In our definition of the ideal-world, either both the adversary and the honest party obtain their output or neither does. This, however, is not necessarily the case in the real-world, where it might be possible for the adversary to abort after having received its own output, but before the honest party obtains its output.

It is a classical result that without an honest majority, it cannot be guaranteed that this cannot happen [Cle86], and requiring an honest majority in a two-party setting is equivalent to requiring that both parties are honest. To fix this, we specify that in the ideal-world, the trusted party should first send the output to Sim, who then gets to decide if the trusted party should also send the honest party its output, but we will generally omit this detail, and just let it be implicitly understood that if the adversary aborts before an honest party would obtain its output, the simulator should abort the ideal functionality. In the literature this is often refereed to as "security with abort".

The final part of our definition of security is the notion of a *hybrid model*. Often we will describe a protocol, Π, for a functionality, $\mathcal{F}$, using another functionality, $\mathcal{G}$, as a subroutine. Here we will say that Π is in the $\mathcal{G}$-hybrid model. When proving the security of a protocol in the $\mathcal{G}$-hybrid model we note that our ideal-world simulator Sim will be simulating the ideal functionality for $\mathcal{G}$, when interacting with $\mathcal{A}$ as a subroutine. Hence, it sees all input to the $\mathcal{G}$ and provides all output from $\mathcal{G}$. The following composition theorem shows why proving security in the hybrid model is still interesting.

Theorem 2.11 *Composition Theorem.*
Let $\mathcal{F}$ and $\mathcal{G}$ be functionalities. Assume that Π implements $\mathcal{F}$ with computational security in the $\mathcal{G}$-hybrid model, and that Γ implements $\mathcal{G}$ with computational security. Then $\Pi^{\mathcal{G}/\Gamma}$ implements $\mathcal{F}$ with computational secu-

rity.

Here $\Pi^{\mathcal{G}/\Gamma}$ denotes the composed protocol, i.e. a protocol like Π, but where the calls to $\mathcal{G}$ has been replaced with an execution of Γ. We have illustrated this concept in Figure 3.

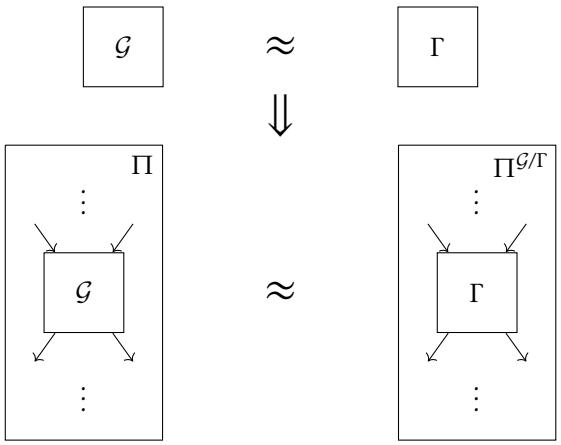

Figure 3: Illustration of the composition theorem.

2.2.2 Examples of functionalities

We will now informally present some examples of functionalities, which we will use or implement later on. The first one is $\mathcal{F}_{\mathsf{ot}}$, the functionality that will later create a protocol that implements.

Example 2.12 $\mathcal{F}_{\mathsf{ot}}$.
We refer to the parties A and B as the Sender S and the Receiver R. The functionality is as follows. S inputs two messages, m_0 and m_1, into $\mathcal{F}_{\mathsf{ot}}$, and R inputs a choice bit $c \in \{0,1\}$ into $\mathcal{F}_{\mathsf{ot}}$. The functionality then outputs m_c to R, and the Sender has no output.

This is often refereed to as $\binom{2}{1}$-OT, and other variants of oblivious transfer exists, eg. $\binom{n}{k}$-OT (Sender inputs n messages, and Receiver chooses $k < n$ messages), and rabin-OT (Sender inputs one message, the Receiver receives the message (alternatively $\perp$) with probability 1/2).

If one now was to go back to Theorem 2.8, one could now restate it as "For every multi-party functionality F, there exists a protocol that implements F with computational security in the $\mathcal{F}_{\mathsf{OT}}$-hybrid model."

Turning our attention to commitments, we define two versions, a typical bit-commitment functionality, and a slightly more general functionality, called selective-opening commitments.

Example 2.13 $\mathcal{F}_{\mathsf{com}}$.
Alice has a bit $a \in \{0,1\}$, Bob has no input. The functionality has two phases,

first the commitment phase where Alice sends a to $\mathcal{F}_{\mathsf{com}}$ and $\mathcal{F}_{\mathsf{com}}$ then sends an acknowledgment for having received a value to Bob. Later, in the opening phase, Alice can tell $\mathcal{F}_{\mathsf{com}}$ to open a to Bob, in which case $\mathcal{F}_{\mathsf{com}}$ sends a to Bob. Otherwise Bob has no output, and in either case Alice has non.

Example 2.14 $\mathcal{F}_{\mathsf{so-com}}$.
Alice has a string $m \in \{0,1\}^n$, Bob has $t \subseteq [n]$. In the commitment phase Alice sends m to $\mathcal{F}_{\mathsf{so-com}}$ and $\mathcal{F}_{\mathsf{so-com}}$ then sends an acknowledgment for having received its input to Bob. In the opening phase Bob sends t to $\mathcal{F}_{\mathsf{so-com}}$, which forwards t to Alice, and Alice can then permit that $\mathcal{F}_{\mathsf{so-com}}$ opens m_t to Bob. Otherwise Bob has no output, and in either case Alice has non.

This last functionality can be seen as equivalent to running n instances of $\mathcal{F}_{\mathsf{com}}$ in parallel. However, when creating protocols for these functionalities it is often simpler to prove sequential security, rather than parallel security, so we will use $\mathcal{F}_{\mathsf{so-com}}$ rather than $\mathcal{F}_{\mathsf{com}}$.

Remark 2.15
When using $\mathcal{F}_{\mathsf{so-com}}$ we will often not work with strings in $\{0,1\}^n$, but with sequences of the form $\{x_i, \theta_i\}_{i\in[n]}$ where $x_i \in \{0,1\}$ and $\theta_i \in \{+,\times\}$. However, mapping $+ \mapsto 0; \times \mapsto 1$ and $\{x_i, \theta_i\} \mapsto \{x_{2i-1}, \theta_{2i}\}$ allows us to use the string-version of $\mathcal{F}_{\mathsf{so-com}}$ with $t \mapsto \{2i-1 : i \in t\} \cup \{2i : i \in t\}$. Hence, we will just refer to $\mathcal{F}_{\mathsf{so-com}}$ when working with sequences on this form.

In practise, our protocols will often be using $\mathcal{F}_{\mathsf{so-com}}$ as a subroutine, where if Alice rejects revealing m_t to Bob, then Bob aborts. In these cases we will omit noting that $\mathcal{F}_{\mathsf{so-com}}$ forwards t to Alice and that Alice permits opening m_t, and just let it be implicitly understood that if Alice rejects, then Bob aborts.

2.3 Sampling games in a classical setting

This section consist of a presentation of sampling games in a classical setting. This serves to both introduce the concept, and we obtain some results that we will see are also meaningful in the quantum setting. This presentation is based on the third section of Bouman and Fehr's paper on quantum sampling [BF10]. Let $q = (q_1, q_2, \ldots, q_n) \in \{0,1\}^n$ be a string of fixed length n and let t be the indexes of a subset of q, i.e. $t \subseteq [n]$. We wish to estimate the relative Hamming distance between the remaining string $q_{\bar{t}}$ and some fixed reference string $q_\circ \in \{0,1\}^n$ (i.e. $\omega_{q_\circ}(q_{\bar{t}}) = \omega((q \oplus q_\circ)_{\bar{t}})$), by only looking at the elements in q_t. To simplify notation, we consider the relative Hamming weight ($\omega(q_{\bar{t}})$, i.e. the relative distance to the zero string) instead of the relative Hamming distance to some reference string, but note that for any reference string $q_\circ$ and string q we have that $\omega_{q_\circ}(q) = \omega(q \oplus q_\circ)$, so any results obtained for estimating the relative Hamming weight are directly applicable to estimating the relative Hamming distance.

The typical way to do this, would be to let t be some uniformly random subset of $[n]$ of some fixed size k, and use $\omega(q_t)$ as our estimate. However, more generally

we will allow any strategy that picks $t \subseteq [n]$ according to some probability distribution P_T, and computes the estimate for $\omega(q_{\bar{t}})$ as some (randomised) function f of t, q_t and a seed s that is sampled according to some probability distribution P_S. This motivates the following definition.

Definition 2.16 *Sampling strategy.*
A sampling strategy Ψ is a triple (P_T, P_S, f) where P_T is a distribution over subsets of $[n]$, P_S is an independent distribution over a finite set $\mathcal{S}$, and f is a function

$$f\colon \{(t,v)\colon t \subseteq [n], v \in \{0,1\}^{|t|}\} \times \mathcal{S} \to \mathbb{R}.$$

Note that since it will make the definition conceptually easier to work with if $P_{TS} = P_T P_S$, we require the seed s to be chosen independently of the subset t. However, it will sometimes be convenient to allow the seed to depend on t. To make this compliant with our definition we use a "container seed" that contains a seed for every possible t, each chosen with the "correct" distribution, and as part of the function f we then extract the seed corresponding to t from the container seed. For an example where we use this see Example 2.21.

2.3.1 The error of a sampling strategy

Naturally, we will be interested in "how well" a sampling strategy performs. Formally, we define "how well" in terms of the probability that our estimate $f(t, q_t, s)$ is how close to the real value $\omega(q_{\bar{t}})$. For a sampling strategy $\Psi = (P_T, P_S, f)$, we introduce the following notation. For arbitrary but fixed $t \in T, s \in S$ and $\delta > 0$ we define

$$B^{\delta}_{t,s}(\Psi) := \{b \in \{0,1\}^n\colon |\omega(b_{\bar{t}}) - f(t, b_t, s)| < \delta\} \subseteq \{0,1\}^n, \tag{2.1}$$

which one can conceptually think of as the "ball" of strings in $\{0,1\}^n$, where our estimate is within δ of the real value, parameterised by t and s. Note that what this definition implies is that if we know that a string b is in $B^{\delta}_{t,s}$, then we are guaranteed that $f(t, b_t, s)$ is δ-close to the real value of $\omega(b_{\bar{t}})$. When Ψ is clear from the context we will write $B^{\delta}_{t,s}$ instead of $B^{\delta}_{t,s}(\Psi)$.

Replacing s and t with the corresponding random variables T and S, we obtain $B^{\delta}_{T,S}$, which is a function of random variables, and hence itself a random variable, with a range consisting of subsets of $\{0,1\}^n$. Using this random variable, we can now define the error probability of a sampling strategy as follows.

Definition 2.17 *Error probability.*
The classical error probability of a sampling strategy $\Psi = (P_T, P_S, f)$ is parameterised by $0 < \delta < 1$ and defined as

$$\varepsilon^{\delta}_{\text{class}}(\Psi) := \max_{q \in \{0,1\}^n} \Pr\left[q \notin B^{\delta}_{T,S}(\Psi)\right]. \tag{2.2}$$

That is, the maximal probability that the estimate of a string in $\{0,1\}^n$ is not δ-close to the real value. Hence, this definition guarantees that for any string $q \in \{0,1\}^n$, the estimate is δ-close to the real value, except with probability at most $\varepsilon^{\delta}_{\text{class}}(\Psi)$. When Ψ is clear from the context we write $\varepsilon^{\delta}_{\text{class}}$ instead of $\varepsilon^{\delta}_{\text{class}}(\Psi)$.

Remark 2.18 *Using a sampling strategy for a test.*
Note that a sampling strategy can be used as a test for closeness of $q_{\overline{t}}$ to the zero string, by accepting if $f(t, q_t, s) = 0$ and rejecting if $f(t, q_t, s) \neq 0$. In this case $\varepsilon^{\delta}_{\text{class}}$ is the worst case probability of accepting, despite $q_{\overline{t}}$ not being δ-close to the zero string, i.e. that $\omega(q_{\overline{t}}) > \delta$.

2.3.2 Examples of sampling strategies

In this section we introduce and analyse a selection of sampling strategies, adding details to the work done by Bouman and Fehr, and combine their examples with the relevant appendices [BF10]. Of special interest is Example 2.21, which we will later make use of.

Example 2.19 *Random sampling without replacement.*
For a parameter k, one uniformly at random chooses k distinct indices $i_1, \ldots, i_k$ within $[n]$, and the relative Hamming weight of $q_{\{i_1,\ldots,i_k\}}$ is used as an estimate for the relative Hamming weight of $q_{\overline{\{i_1,\ldots,i_k\}}}$. Hence, the sampling strategy is given by $\Psi_1 = (P_T, P_S, f)$ where $P_T(t) = 1/\binom{n}{k}$ if $|t| = k$ and $P(t) = 0$ otherwise, $\mathcal{S} = \{\perp\}$ with $P_S(\perp) = 1$, and $f(t, q_t, s) = \omega(q_t)$.

Towards finding $\varepsilon^{\delta}_{\text{class}}$, we first note that considering the k bits sampled as random variables $T_1, \ldots, T_k$, we have that by definition $\omega(q_T) = \frac{1}{k}\sum T_i$. Hence, it follows immediately from Hoeffding's inequality (Theorem 2.5) that the estimate is δ-close to the relative Hamming weight of q, except with probability at most $2\exp(-2\delta^2 k)$.

To instead estimate $\omega(q_{\overline{T}})$, we note that by letting $\alpha = k/n$, we can write $\omega(q)$ as $\omega(q) = \alpha\omega(q_T) + (1-\alpha)\omega(q_{\overline{T}})$. It follows that

$$\begin{aligned}
\omega(q_{\overline{T}}) - \omega(q_T) &= \frac{1}{1-\alpha}(\omega(q) - \alpha\omega(q_T)) - \omega(q_T) \\
&= \frac{1}{1-\alpha}(\omega(q) - \alpha\omega(q_T) - (1-\alpha)\omega(q_T)) \\
&= \frac{1}{1-\alpha}(\omega(q) - \omega(q_T)), \qquad (2.3)
\end{aligned}$$

and therefore

$$\begin{aligned}
\varepsilon^{\delta}_{\text{class}} &\overset{(2.2)}{=} \max_{q\in\{0,1\}^n} \Pr\left[q \notin B^{\delta}_{T,S}\right] \\
&\overset{(2.1)}{=} \max_{q\in\{0,1\}^n} \Pr\left[|\omega(q_{\overline{T}}) - \omega(q_T)| \geq \delta\right] \\
&\overset{(2.3)}{=} \max_{q\in\{0,1\}^n} \Pr\left[\left|\frac{1}{1-\alpha}(\omega(q) - \omega(q_T))\right| \geq \delta\right]
\end{aligned}$$

$$= \max_{q\in\{0,1\}^n} \Pr\left[|\omega(q) - \omega(q_T)| \geq (1-\alpha)\delta\right]$$
$$\leq 2\exp\left(-2(1-\alpha)^2\delta^2 k\right), \tag{2.4}$$

where the last line is once again Hoeffding's inequality. If we additionally assume that $k \leq n/2$, we obtain the following simpler upper bound

$$\varepsilon^{\delta}_{\text{class}} \overset{(2.4)}{\leq} 2\exp\left(-2(1-1/2)^2\delta^2 k\right) = 2\exp\left(-\frac{1}{2}\delta^2 k\right). \tag{2.5}$$

Example 2.20 *Uniformly random subset sampling.*
One uniformly at random chose any $t \in \mathcal{P}([n])$, and the estimate for the relative Hamming weight is simply the relative Hamming weight of q_t.[2] Hence, the sampling strategy is given by $\Psi_2 = (P_T, P_S, f)$ where $P_T(t) = 1/2^n$ for every $t \subseteq [n]$, $\mathcal{S} = \{\perp\}$ with $P_S(\perp) = 1$, and $f(t, q_t, s) = \omega(q_t)$.

Towards finding $\varepsilon^{\delta}_{\text{class}}$ we first note that for any fixed size k of t, t is obtained as in Example 2.19, and hence we have a bound for the error in this case.

Now, we can think of picking t as drawing n uniformly random bits, where drawing a 1 indicates that we include the entry in t and a 0 indicates that we do not. Letting X_i be the random variable for the i'th bit drawn we see that

$$\overline{X} = \frac{1}{n}\sum_{i=1}^{n} X_i = \frac{k}{n},$$

where $k = |t|$. Hence, Hoeffding's inequality implies that for $0 < \beta < \frac{1}{2}$ we have[3]

$$\Pr\left[\left|\overline{X} - \frac{1}{2}\right| \geq \beta\right] \leq 2\exp\left(-2\beta^2 n\right).$$

We now assume that $k \in \left[\left(\frac{1}{2}-\beta\right)n, \left(\frac{1}{2}+\beta\right)n\right]$, and include the probability of this not being the case in our final error. As it was previously noted, for any fixed k Equation (2.4) holds, i.e.

$$\varepsilon^{\delta,k}_{\text{class}} \leq 2\exp\left(-2\left(1-\frac{k}{n}\right)^2\delta^2 k\right).$$

To obtain an upper bound that is independent of k, we replace the first occurrence with an upper bound for k, i.e. $\left(\frac{1}{2}+\beta\right)n$, and the second occurrence with a lower bound, i.e. $\left(\frac{1}{2}-\beta\right)n$. Doing this we obtain that

$$\varepsilon^{\delta,k}_{\text{class}} \leq 2\exp\left(-2\left(1-\frac{\left(\frac{1}{2}+\beta\right)n}{n}\right)^2\delta^2\left(\frac{1}{2}-\beta\right)n\right)$$
$$= 2\exp\left(-2n\delta^2\left(1-\frac{1}{2}-\beta\right)^2\left(\frac{1}{2}-\beta\right)\right)$$

[2] Note that we technically allow that $t = \emptyset$, which would make for a very poor sample size, and that in practise one should probably require a minimum size of t.
[3] Note that we are here using Hoeffding's inequality in the general probabilistic setting, and not the restricted-to-relative-Hamming-weight setting that we are usually using it in [Hoe63].

$$= 2\exp\left(-2n\delta^2\left(\frac{1}{2}-\beta\right)^3\right).$$

We can now find an upper bound for the total error of this sampling strategy by combing the probability that k is not in this interval with the upper bound for the probability that if k is in this interval, then our estimate is more than δ away from the true value,

$$\varepsilon^{\delta}_{\text{class}} \leq 2\exp\left(-2n\delta^2\left(\frac{1}{2}-\beta\right)^3\right) + 2\exp\left(-2\beta^2 n\right).$$

Finally, we can simplify this by letting $\beta = \delta/4$ (which is an allowable β-value), which gives us

$$\varepsilon^{\delta}_{\text{class}} \leq 2\exp\left(\frac{-n\delta^2(2-\delta)^3}{32}\right) + 2\exp\left(\frac{-n\delta}{8}\right),$$

and since $0 < \delta < 1$ an upper bound for both exponents is $\frac{-n\delta^2}{32}$, so our final (simpler) bound is

$$\varepsilon^{\delta}_{\text{class}} \leq 4\exp\left(\frac{-n\delta^2}{32}\right). \tag{2.6}$$

Example 2.21 *Random sampling without replacement, using only part of the sample.*
The final example we will consider might at first appear unnatural, but as we will later see, this is exactly the situation that we will end up in. It is in some sense a combination of Example 2.19 and 2.20, i.e. first we choose a random $t \subseteq [n]$ of a fixed size k, as in Example 2.19 such that $P_T(t) = 1/\binom{n}{k}$ if $|t| = k$ and $P(t) = 0$ otherwise. However, we only use a uniformly randomly chosen subset $s \subseteq t$ to compute the estimate, i.e.

$$f(t, q_t, s) = \omega(q_s),$$

with $P_S(s) = 1/2^k$ for any $s \subseteq t$. Note that while this description of P_S implicitly depends on the choice of t, we consider the version of it that is constructed as was suggested after Definition 2.16. Thus, P_S is independent of P_T, and our sampling strategy is consistent with Definition 2.16.

To find the error probability of this sampling strategy, we note that we know from Equation (2.5) of Example 2.19 that for $k < n/2$ we have

$$\Pr\left[\left|\omega(q_{\overline{T}}) - \omega(q_T)\right| \geq \xi\right] \leq 2\exp\left(-\frac{1}{2}\xi^2 k\right).$$

Further, we can consider the selection of the seed s as a uniformly random subset sampling of q_t, so it follows from Equation (2.6) of Example 2.20 that

$$\Pr\left[|\omega(q_T) - \omega(q_S)| \geq \gamma\right] \leq 4\exp\left(-\frac{1}{32}\gamma^2 k\right).$$

Finally, letting $\delta = \xi + \gamma$ we obtain that

$$\begin{aligned}
\varepsilon^{\delta}_{\text{class}} &= \max_{q\in\{0,1\}^n} \Pr\left[\left|\omega(q_S) - \omega(q_{\overline{T}})\right| \geq \delta\right] \\
&\leq \max_{q\in\{0,1\}^n} \Pr\left[\left|\omega(q_S) - \omega(q_T)\right| + \left|\omega(q_T) - \omega(q_{\overline{T}})\right| \geq \gamma + \xi\right] \\
&\leq \min_{\substack{0<\xi,\gamma<\delta \\ \xi+\gamma=\delta}} \left[4\exp\left(-\frac{1}{32}\gamma^2 k\right) + 2\exp\left(-\frac{1}{2}\xi^2 k\right)\right] \qquad (2.7) \\
&= \min_{0<\xi<\delta} \left[4\exp\left(-\frac{1}{32}(\delta-\xi)^2 k\right) + 2\exp\left(-\frac{1}{2}\xi^2 k\right)\right] \\
&\leq 6\exp\left(-\frac{1}{50}k\delta^2\right), \qquad (2.8)
\end{aligned}$$

where (2.7) follows from Boole's inequality, and (2.8) follows from setting $\xi = \frac{\delta}{5}$, so that the exponents agree.

3 Introducing the quantum world

In Section 4 and 5 we will introduce quantum cryptography, and see how it differs from classical cryptography. Before we can do so we have to introduce quantum computation theory and quantum information theory. This section is loosely based on the classical book by Nielsen and Chuang [NC02], except when stated otherwise.

Quantum information theory is reminiscent of classical information theory, but adapted to take into account the effects of quantum mechanics. We present this theory in Section 3.3, and show a number of results. We note that the term "quantum information" is sometimes used as a broader catch-all term, relating to everything that can be understood as relating to information processing using quantum mechanics, eg. also quantum computation. To avoid confusion we will always use the term "quantum information theory" when referring to the study of the quantum analogy to classical information theory.

So what about quantum computation theory? In a nutshell, a quantum computation is a computation where we (ab)use quantum mechanics, and hence quantum computation theory is the study of how the classical topics of computation theory change, when one considers quantum mechanics. This is the focus of Section 3.2.

Seeing that both quantum information theory and quantum computation theory relies on quantum mechanics, it is natural to take a step further back, and start with an introduction to quantum mechanics.

3.1 Quantum mechanics

Quantum mechanics is a mathematical framework for constructing physical theories, such as quantum electrodynamics. The rules of quantum mechanics are simple, yet counterintuitive. Fundamentally, it is these rules that quantum computation- and information theory has to adhere to.

There are multiple ways to define quantum mechanics. We follow Nielsen and Chuang [NC02], using four postulates that provides the connection between the physical world and our mathematical formalism. We present the four postulates, and show how they imply the notation we use. In the last subsections we show some additional results, and introduce some convenient notation.

3.1.1 Postulate 1: States

> **Postulate 1**
> Associated to any isolated physical system is a complex Hilbert space $\mathcal{H}$ known as the state space of the system. The system is completely described by its state vector $|\phi\rangle$, which is a unit vector in the system's state space.

While the four postulates will define quantum mechanics, there is still a lot of freedom when constructing a specific theory. For instance, Postulate 1

say that the state space is a Hilbert space, but not *what* Hilbert space it is. The choice is dependent on the physical system one is modelling. For instance, for modelling the position and momentum states of a single non-relativistic spin zero particle it is common to use $L^2(\mathbb{C})$. Luckily, we will only be interested in a much simpler system, namely that of a *qubit*. A qubit is the quantum analogy of a classical bit, and has $\mathbb{C}^2$ as its state spaces. Hence, the state vector of a qubit is a unit vector in $\mathbb{C}^2$.

Remark 3.1 *Bases for qubits.*
The state vector of a qubit can be written in any basis for $\mathbb{C}^2$, but we will mainly be interested in the computational basis $\{|0\rangle, |1\rangle\}$, and the Hadamard basis $\{|+\rangle, |-\rangle\}$. We refer to the computational basis as "+" and the Hadamard basis as "×". The basis vectors are defined as

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad |+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}, \quad |-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}} = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{-1}{\sqrt{2}} \end{pmatrix}.$$

We can then write the state of a qubit $|\phi\rangle$ as

$$|\phi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{3.1}$$

and similarly $\alpha|0\rangle + \beta|1\rangle$ is a valid state of a qubit if $|\alpha|^2 + |\beta|^2 = 1$. We will then say that $|\phi\rangle$ is in a superposition of states $|0\rangle$ and $|1\rangle$, with amplitude α and β, respectively.

This first example already teases the power that quantum has. Where a classical bit can only take on one out of two values, a qubit can be in an infinite number of superpositions.

3.1.2 Postulate 2: Transformation

Postulate 2
The evolution of a closed quantum system is described by a unitary transformation U.

Postulate 2 say that the discrete evolution of a state from a time t to a later time t', can be described by a unitary transformation. It does not, however, say when a unitary transformation describes the evolution of a system. Luckily for us, it turns out that any unitary can be realised as the transformation of a system. Some examples of unitary operators are the Pauli matrices

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Another typical example is the Hadamard gate

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix},$$

which changes between the computational basis and the Hadamard basis, in the sense that $H|0\rangle = |+\rangle, H|1\rangle = |-\rangle, H|+\rangle = |0\rangle$, and $H|-\rangle = |1\rangle$.

3.1.3 Postulate 3: Measurements

Postulate 3
Quantum measurements are described by a collection of measurement operators $\{M_m\}$. These are operators acting on the state space of the system being measured. The index m refers to the measurement outcomes that may occur. If the state of the system is $|\phi\rangle \in \mathcal{H}$ immediately before measuring, then the probability of observing outcome m is given by

$$p(m) = \langle\phi|M_m^* M_m|\phi\rangle,$$

and the state of the system after observing outcome m is

$$\frac{M_m|\phi\rangle}{\sqrt{\langle\phi|M_m^* M_m|\phi\rangle}}.$$

The measurement operators satisfy the completeness equation

$$\sum_m M_m^* M_m = \mathbb{I}.$$

While postulate 2 say that the evolution of a *closed* system can be described by a unitary, there must come a time when the outside world interacts with the system, namely when we observe it (otherwise the system is of very questionable interest). At this point, the system is no longer closed, and hence it is not necessarily subject to only unitary operations. This is where Postulate 3 comes in, describing the effects of measuring the system.

Example 3.2
Maybe the simplest, but non the less important, example is measuring the state of a qubit with respect to the computational basis. Consider the qubit $|\phi\rangle$ from (3.1). We have the measurement operators $M_0 := |0\rangle\langle 0|$ and $M_1 := |1\rangle\langle 1|$, and observe that

$$M_0 + M_1 = |0\rangle\langle 0| + |1\rangle\langle 1| = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \mathbb{I}.$$

The probability of observing 0 is then

$$p(0) = \langle\phi|M_0^* M_0|\phi\rangle = \langle\phi|M_0|\phi\rangle = |\alpha|^2,$$

and similarly, observing 1 happens with probability $|\beta|^2$. Assuming that we observe 0, the state after measuring is

$$\frac{M_0|\phi\rangle}{\sqrt{\langle\phi|M_0^* M_0|\phi\rangle}} = \frac{\alpha}{|\alpha|}|0\rangle.$$

Example 3.3
Let us further refine the above example. Assume that $|\phi\rangle$ that is either $|+\rangle$ or $|-\rangle$. Upon measuring with respect to M_0, M_1 we obtain that for $|+\rangle$ we have

$$p(0) = \langle +|M_0^* M_0|+\rangle = \frac{1}{2} = \langle +|M_1^* M_1|+\rangle = p(1),$$

and similarly for $|-\rangle$. Additionally, it is clear that if we observe 0 then both $|+\rangle$ and $|-\rangle$'s post-measurement state would be $|0\rangle$. If we observe 1, however, the post-measurement states are $|1\rangle$ and $-|1\rangle$, respectively.

If we now invoke Theorem 3.4 (which we will show right after finishing this example), and we conclude that if we measure $|+\rangle$ or $|-\rangle$ with respect to the computational basis, we obtain 0 and 1 with equal probability, and our post-measurement state is the same, regardless of which one we started with.

> **Theorem 3.4** *Global phase.*
> Let $|\phi\rangle$ be a state vector. For any $\theta \in \mathbb{R}$ we say that $|\phi\rangle$ is equal to the state vector $e^{i\theta}|\phi\rangle$ up to a global phase factor $e^{i\theta}$.
>
> We claim that the statistics of measuring $|\phi\rangle$ and $e^{i\theta}|\phi\rangle$ are the same, and hence we can ignore the global phase.

Proof. We show that for arbitrary measurements one obtain the same outcomes from measuring $|\phi\rangle$ and $e^{i\theta}|\phi\rangle$. Let M_m be any measurement operator associated to some quantum measurement. Then

$$\langle\phi|e^{-i\theta} M_m^* M_m e^{i\theta}|\phi\rangle = \langle\phi|e^0 M_m^* M_m|\phi\rangle = \langle\phi|M_m^* M_m|\phi\rangle,$$

so the measurement probabilities are the same for both states. □

The measurements M_0 and M_1, that we considered in Example 3.2 and 3.3, are the simplest form of measurements, and Postulate 3 allow many other collections of measurement operators. The two families of measurement collections that one typically studies is projective measurements, and positive operator-valued measurements (POVMs). We present these two families in Example 3.5 and 3.6.

Example 3.5 *Projective measurements.*
A measurement collection $\{M_m\}$ is said to consist of projective measurements, if the measurement operators are orthogonal projections.

A common example of a projective measurement is the measurements derived from a basis. If $\{|i\rangle\}_{i\in I}$ is an orthonormal basis for a Hilbert space, then it is easily verifiable that $\{|i\rangle\langle i|\}_{i\in I}$ consists of orthogonal projections on the Hilbert space, and that

$$\sum_{i\in I} |i\rangle\langle i| = \mathbb{I}.$$

Thus, the measurements that we considered in Example 3.2 and 3.3 was in fact projective measurements.

Example 3.6 *POVMs.*
Often times, one is not interested in the post-measurement state, but only in the probabilities of the different measurement outcomes. In this case, the POVM formalism can be very useful. Suppose that $\{E_m\}$ is a set of positive operators such that $\sum_m E_m = \mathbb{I}$. Since each E_m is positive, $\sqrt{E_m}$ is well-defined, and noting that $\{\sqrt{E_m}\}$ forms a measurement collection, in the sense of Postulate 3, it is clear that $\{E_m\}$ is all one need if one is only interested in the outcomes, since

$$\langle\phi|(\sqrt{E_m})^*\sqrt{E_m}|\phi\rangle = \langle\phi|E_m|\phi\rangle.$$

If one just wanted to do physics, considering these two families would generally be sufficient. However, at a later time we will want to argue about distinguishing states, and to be able to do so with generality, we need to consider general measurements, and not just select families. The following theorem indicates why this is needed, showing how distinguishability is different in the quantum world.

Theorem 3.7
Non-orthogonal states cannot be reliably distinguished.

Proof. Assume towards a contradiction that a measurement collection $\{M_m\}$ distinguish the non-orthogonal states $|\phi\rangle$ and $|\psi\rangle$ with certainty, i.e. we have some function f that maps the possible outcomes of measuring to $\{\phi, \psi\}$, such that if the state that was measured was $|\phi\rangle$ it always maps to ϕ, and vice versa. Then the assumption is equivalent to stating that when measuring $|\phi\rangle$, the probability of measuring i such that $f(i) = \phi$ is 1. If one now defines

$$E_\phi := \sum_{\substack{i:\\ f(i)=\phi}} M_i^* M_i, \qquad\qquad E_\psi := \sum_{\substack{i:\\ f(i)=\psi}} M_i^* M_i,$$

one has that

$$\langle\phi|E_\phi|\phi\rangle = 1 = \langle\psi|E_\psi|\psi\rangle.$$

Since $E_\phi + E_\psi = \mathbb{I}$, it must also be the case that $\langle\phi|E_\phi|\phi\rangle + \langle\phi|E_\psi|\phi\rangle = 1$, and hence $\langle\phi|E_\psi|\phi\rangle = 0$, and so is $\sqrt{E_\psi}|\phi\rangle$.

Decomposing $|\psi\rangle$ as $|\psi\rangle = \alpha|\phi\rangle + \alpha|\theta\rangle$ with $|\theta\rangle \perp |\phi\rangle$ we observe that $|\beta| < 1$, since $|\phi\rangle$ and $|\psi\rangle$ are assumed to not be orthogonal and $|\alpha|^2 + |\beta|^2 = 1$. Hence, $\sqrt{E_\psi}|\psi\rangle = \beta\sqrt{E_\psi}|\theta\rangle$, therefore

$$1 = \langle\psi|E_\psi|\psi\rangle = |\beta|^2\langle\theta|E_\psi|\theta\rangle \leq |\beta|^2 < 1,$$

which is a contradiction. □

3.1.4 Postulate 4: Composite systems and entanglement

Postulate 4
The state space of a composite physical system is the tensor product of the state spaces of the components. Moreover, if the systems are labelled 1 through n, and each system is prepared in state $|\phi_i\rangle$, then to joint state of the total system is

$$|\phi_1\rangle \otimes |\phi_2\rangle \otimes \cdots \otimes |\phi_n\rangle.$$

This final postulate indicates how we can combine multiple systems. In fact, we will mostly be working with composite systems of multiple qubits, the quantum equivalent to classical bit-strings.

Example 3.8
The simplest composite system we will work with consists of two qubits. Its state space is then $\mathbb{C}^2 \otimes \mathbb{C}^2 \simeq \mathbb{C}^4$, and a basis for it can be found by taking the tensor products of a basis of $\mathbb{C}^2$;

$$|0\rangle \otimes |0\rangle = |00\rangle = \begin{pmatrix}1\\0\end{pmatrix} \otimes \begin{pmatrix}1\\0\end{pmatrix} = \begin{pmatrix}1\\0\\0\\0\end{pmatrix}, \qquad |0\rangle \otimes |1\rangle = |01\rangle = \begin{pmatrix}1\\0\end{pmatrix} \otimes \begin{pmatrix}0\\1\end{pmatrix} = \begin{pmatrix}0\\1\\0\\0\end{pmatrix},$$

$$|1\rangle \otimes |0\rangle = |10\rangle = \begin{pmatrix}0\\1\end{pmatrix} \otimes \begin{pmatrix}1\\0\end{pmatrix} = \begin{pmatrix}0\\0\\1\\0\end{pmatrix}, \qquad |1\rangle \otimes |1\rangle = |11\rangle = \begin{pmatrix}0\\1\end{pmatrix} \otimes \begin{pmatrix}0\\1\end{pmatrix} = \begin{pmatrix}0\\0\\0\\1\end{pmatrix}.$$

Another concept that Postulate 4 enables us to define and study is entanglement. The importance of entanglement cannot be understated, and it is hypothesised that entanglement might be the feature that set quantum computation apart from classical computation [JL03].

Definition 3.9 *Entanglement.*
A state with the property that it cannot be written as the product of states of its component systems is said to be entangled.

Example 3.10 *EPR-pair.*
Consider the following two-qubit state

$$|\phi\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}.$$

A two-qubit state of this form is called an *EPR-pair*.

There exists no one-qubit states $|a\rangle$ and $|b\rangle$ such that $|\phi\rangle = |a\rangle \otimes |b\rangle$. To see

this, note that such $|a\rangle$ and $|b\rangle$ would have to satisfy that

$$|a\rangle \otimes |b\rangle = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \otimes \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = \begin{pmatrix} a_1 b_2 \\ a_1 b_2 \\ a_2 b_1 \\ a_2 b_2 \end{pmatrix} \stackrel{?}{=} \begin{pmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ 0 \\ \frac{1}{\sqrt{2}} \end{pmatrix} = \frac{|00\rangle + |11\rangle}{\sqrt{2}} = |\phi\rangle,$$

which is clearly unsitisfiable. This shows that $|\phi\rangle$ is entangled.

Assume now that we were to measure the first qubit in the computational basis, and obtain a 0. How would this affect the second qubit? To answer this we note that a measurement collection for measuring the first qubit in the computational basis is given by $\{|0\rangle\langle 0| \otimes \mathbb{I}, |1\rangle\langle 1| \otimes \mathbb{I}\}$. Writing out $|0\rangle\langle 0| \otimes \mathbb{I}$, we see that

$$\begin{aligned} |0\rangle\langle 0| \otimes \mathbb{I} &= \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \otimes \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \end{aligned}$$

so that

$$(|0\rangle\langle 0| \otimes \mathbb{I})|\phi\rangle = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ 0 \\ \frac{1}{\sqrt{2}} \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

It therefore follows that if we obtain a 1, the post-measurement state of $|\phi\rangle$ is

$$\frac{(|0\rangle\langle 0| \otimes \mathbb{I})|\phi\rangle}{\sqrt{\langle\phi|(|0\rangle\langle 0| \otimes \mathbb{I})(|0\rangle\langle 0| \otimes \mathbb{I})|\phi\rangle}} = |00\rangle,$$

so that if we now measure the second qubit in the computational basis, we measure 0 with probability 1.

The last result we derived in Example 3.10 is perhaps one of the most surprising results in quantum mechanics! By interacting with only the first qubit in a two-qubit system, it has somehow been determined what the outcome of measuring the second qubit will be. This result was in fact so surprising and counter-intuitive that three of the founders of quantum mechanics — Einstein, Podolsky, and Rosen — published a paper where they argued that a variation of this "paradox" showed that quantum mechanics was not complete. Hence, this result is known as the EPR-paradox, despite not being a paradox, and having since been verified to hold in practice.

In fact, many famous and practically verified results in quantum mechanics are applications of entanglement. An example of such is the non-local CHSH

game. In this game, two players, Alice and Bob, are allowed to agree on a strategy and communicate, but only before the game starts. Once the game starts, the players are each given a uniformly random bit x, y, and then they each have to output a bit (a and b), and they win if $xy = a \oplus b$. If Alice and Bob uses only classical tools, the optimal strategy gives them a 3/4 chance of winning. However, by sharing an EPR-pair in the pre-game phase, Alice and Bob can win with probability $\cos^2(\pi/8) \approx 0.85$ [Bru+14].

3.1.5 Density operators

So far we have formulated quantum mechanics in a language of state vectors, but an alternative formulation can be made using *density operators* (also called *density matrices*). Both formulations are equivalent, and which one is most practical depends on the application.

Suppose that a quantum system is in one of a number of states $|\phi_i\rangle$, with respective probabilities p_i. We then call $\{p_i, |\phi_i\rangle\}$ an ensemble (of pure states), and define the density operator for the system by

$$\rho = \sum_i p_i |\phi_i\rangle\langle\phi_i|.$$

On the other hand, an operator ρ is the density operator for some ensemble of pure states if it has trace 1 and is a positive operator. We formalise this characterization of density operators in the following theorem.

Theorem 3.11
An operator ρ is the density operator associated to some ensemble, $\{p_i, |\phi_i\rangle\}$, if and only if the trace of ρ is 1 and ρ is a positive operator.

Proof. Suppose $\rho = \sum_i p_i |\phi_i\rangle\langle\phi_i|$ is a density operator. Then

$$\operatorname{tr}\rho = \sum_i p_i \operatorname{tr}(|\phi_i\rangle\langle\phi_i|) = \sum_i p_i = 1,$$

and for any $|\psi\rangle$ in the state space we have that

$$\begin{aligned}\langle\psi|\rho|\psi\rangle &= \sum_i p_i \langle\psi|\phi_i\rangle\langle\phi_i|\psi\rangle \\ &= \sum_i p_i |\langle\psi|\phi_i\rangle|^2 \geq 0,\end{aligned}$$

so ρ has trace 1 and is positive.

Conversely, suppose ρ has these properties. Since ρ is positive it has a spectral decomposition

$$\rho = \sum_j \lambda_j |j\rangle\langle j|,$$

where each $\lambda_j \in \mathbb{R}^+$, and the vectors $|j\rangle$ are orthonormal. Noting that since the sum of the eigenvectors is equal to the trace of ρ, we have that $\sum_i \lambda_i = 1$, and

thus $\{\lambda_i, |i\rangle\}$ is an ensemble of states with associated density operator ρ. □

With this characterization in mind, it is reasonable to define the density operator of a pure state to be a positive operator with trace 1. Doing so, the four postulates of quantum mechanics can be reformulated in terms of density operators, instead of state vectors. We do not do so here, but remark that the equivalent of $U|v\rangle$ in Postulate 2 is $U\rho U^*$, and that in Postulate 3, the probability of observing outcome m is $\mathrm{tr}(M_m^* M_m \rho)$.

Definition 3.12 *Pure and mixed states.*
A quantum system whose state $|\phi\rangle$ is known exactly is said to be a pure state. In this case the density operator is exactly $\rho = |\phi\rangle\langle\phi|$. If a state is not pure then it is said to be a mixed state, and ρ is said to be a mixture of the different pure states in the associated ensemble.

One of the advantages to using the density operator formalism is that it is easy to decide if a state is pure.

Theorem 3.13
Let ρ be a density operator. Then $\mathrm{tr}\,\rho \leq 1$, with equality if and only if ρ is pure.

Proof. There exists an orthogonal ensemble $\{p_i, |i\rangle\}$ such that $\rho = \sum_i p_i |i\rangle\langle i|$. Then

$$\begin{aligned}\rho^2 &= \left(\sum_i p_i |i\rangle\langle i|\right) \\ &= \sum_{i,j} p_i p_j |i\rangle\langle i|j\rangle\langle j|, \\ &= \sum_i p_i^2 |i\rangle\langle i|,\end{aligned}$$

since $\langle i|j\rangle = 0$ when $i \neq j$. Hence, it follows that

$$\begin{aligned}\mathrm{tr}\,\rho^2 &= \mathrm{tr}\left(\sum_i p_i^2 |i\rangle\langle i|\right) \\ &= \sum_i p_i^2 \,\mathrm{tr}\,|i\rangle\langle i| \\ &= \sum_i p_i^2,\end{aligned}$$

and since each $p_i \leq 1$ and $\sum p_i = 1$, the claim follows. □

Another application of the density operator is as a descriptive tool for subsystems of composite systems. Suppose we have physical systems A and B, with joint state ρ_{AB}. The reduced density operator for system A is then

$$\rho_A = \mathrm{tr}_B\, \rho_{AB}.$$

Example 3.14
To see that this definition is sensible consider first the case where A and B are not entangled, i.e. there exists $\sigma \in \mathcal{H}_A$ and $\pi \in \mathcal{H}_B$ such that $\rho_{AB} = \sigma \otimes \pi$. Then we obtain that

$$\rho_A = \mathrm{tr}_B \rho_{AB} = \mathrm{tr}_B (\sigma \otimes \pi) = \sigma \,\mathrm{tr}\, \pi = \sigma,$$

which is the result we intuitively expect. A more interesting result is obtained if we instead consider an EPR-pair, which has density operator

$$\begin{aligned}\rho &= \left(\frac{|00\rangle + |11\rangle}{\sqrt{2}}\right)\left(\frac{\langle 00| + \langle 11|}{\sqrt{2}}\right) \\ &= \frac{|00\rangle\langle 00| + |11\rangle\langle 00| + |00\rangle\langle 11| + |11\rangle\langle 11|}{2}.\end{aligned}$$

We can now find the reduced density operator on the first qubit to be

$$\begin{aligned}\rho_1 &= \mathrm{tr}_2 \rho \\ &= \frac{\mathrm{tr}_2 |00\rangle\langle 00| + \mathrm{tr}_2 |11\rangle\langle 00| + \mathrm{tr}_2 |00\rangle\langle 11| + \mathrm{tr}_2 |11\rangle\langle 11|}{2} \\ &= \frac{|0\rangle\langle 0|\langle 0|0\rangle + |1\rangle\langle 0|\langle 0|1\rangle + |0\rangle\langle 1|\langle 1|0\rangle + |1\rangle\langle 1|\langle 1|1\rangle}{2} \\ &= \frac{1}{2}\mathbb{I},\end{aligned}$$

which is a mixed state, despite the joint state in the two-qubit system being pure. It is straight forward to show that if ρ' was instead generated from $\frac{|01\rangle+|10\rangle}{\sqrt{2}}$ one would obtain the same reduced density operator for the first qubit. One way to interpret this result is that someone holding only the first qubit will "see" $\frac{1}{2}\mathbb{I}$ in either case, and is therefore unable to distinguish the two situations. This concept will be useful for quantum cryptography.

Remark 3.15
The state $\frac{1}{2}\mathbb{I}$ is a special state, and is called the *maximally mixed state*. It has the property that regardless of what orthogonal basis it is measured with respect to, each outcome is equally likely to occur.

This notion generalises to higher dimensional spaces, and for a state space of dimension d, the maximally mixed state is $\frac{1}{d}\mathbb{I}$.

The final tool we introduce in this section is purification. Essentially, this allows us to consider a pure state in a state space of higher dimension, in place of a mixed state.

Theorem 3.16 *Purification.*
For every mixed state ρ_A of a system A, there exists another system R such that one can define a pure state $|AR\rangle$ for the joint state of A and R with the

property that

$$\rho_A = \mathrm{tr}_R(|AR\rangle\langle AR|).$$

Proof. Fix ρ_A, and let $\sum_i p_i |i\rangle_A \langle i|_A$ be an orthonormal decomposition of ρ_A. To purify ρ_A we let R be a system with the same state space as A, and fix an orthonormal basis $\{|i\rangle_R\}$. We then define the pure state

$$|AR\rangle = \sum_i \sqrt{p_i}\,|i\rangle_A\,|i\rangle_R\,,$$

and note that this has the desired property that

$$\begin{aligned}\mathrm{tr}_R(|AR\rangle\langle AR|) &= \sum_{i,j} \sqrt{p_i}\sqrt{p_j}\,|i\rangle_A\langle i|_A \,\mathrm{tr}\,|i\rangle_R\langle i|_R \\ &= \sum_{i,j} \langle i|j\rangle_R \sqrt{p_i}\sqrt{p_j}\,|i\rangle_A\langle i|_A \\ &= \sum_i p_i\,|i\rangle_A\langle i|_A \\ &= \rho_A,\end{aligned}$$

since $\langle i|j\rangle_R = 0$ when $i \neq j$. □

3.1.6 Classical-quantum systems

A special type of composite systems that will be particular interest is classical-quantum systems. This is composite systems where one of the states is completely classical, and the other is quantum. One way these can arise is during the execution of a cryptographic protocol, where at the end Alice has some classical information, about which an adversary holds some quantum information. To describe this structure, we will first see how we can write a classical probability distribution as a density matrix.

Suppose that we have a classical random variable, with distribution P over a set $S = \{1, 2, \ldots, d\}$. Then we can identify the elements in S with basis elements from the standard basis for $\mathbb{C}^{|S|}$, and hence express a source preparing each of these possible states $|s\rangle$ with probability p_s by the mixture

$$\rho_X = \sum_{s\in S} p_s\,|s\rangle\langle s|.$$

Note that this is just a diagonal matrix, with the probabilities on the diagonal. Assume now that depending on the outcome of P, one prepares a quantum state E in one of $|S|$ ways, with associated density operators $\{\rho_s\}_{s\in S}$. Then the joint state of the classical system and the quantum system is given by

$$\rho_{XE} = \sum_{s\in S} p_s\,|s\rangle\langle s| \otimes \rho_s.$$

In general we will call a system with density operator of this form a classical-quantum system, even when it is obtained in a different way.

3.2 Quantum computation theory

Much like the classical theory of computation, quantum computation theory deals with what problems it is possible to solve for a theoretic model of a quantum computer (computability theory), and how fast it is possible to solve these problems (complexity theory). In this project we generally do not go into details on these topics, and hence we also keep this section brief. We note that this is despite the fact that much like their classical counterparts, both topics are highly relevant for cryptography. The first part of this section follows [Hid19].

A good place to start our considerations might be in the classical world. In 1936, Allan Turing invented the Turing Machine, which the Chruch-Turing Thesis tells us can be used to decide if a classical computer can solve a classical problem or not.

> **Theorem 3.17** *Church-Turing Thesis.*
> If an algorithm can run on a piece of hardware, then there is an equivalent algorithm for the Universal Turing Machine (UTM).

Using this thesis, one needs only to consider the UTM to decide if a problem can or cannot be solved by a computer. The UTM is a well-defined mathematical model for computation, and it is therefore much more straight forward to argue about. Since it is straight forward to (inefficiently) simulate a quantum computer on a classical computer, the Church-Turing Thesis is also applicable to deciding if a problem *can* be solved with a quantum computer.

The original Church-Turing Thesis say nothing about how efficient a problem can be solved, and an updated version was therefore introduced, the Strong Church-Turing Thesis. This thesis, however, proved to be insufficient for probabilistic algorithms, and hence one uses the Extended Church-Turing Thesis.

> **Theorem 3.18** *Extended Church-Turing Thesis.*
> Any algorithmic process can be simulated efficiently using a Probabilistic Turing Machine (PTM).

Which still hold for classical computers. However, the thesis is not strong enough for quantum computation, since there are problems that a quantum computer can solve exponentially faster than a classical PTM can. To circumvent this, the Quantum Extended Church-Turing Thesis (QECTT) was introduced. This thesis still stands.

> **Theorem 3.19** *Quantum Extended Church-Turing Thesis.*
> Any realistic physical computing device can be efficiently simulated by a fault-tolerant quantum computer.

With the QECTT one can start classifying the problems that can be solved by a quantum computer, by how efficiently they can be solved, extending the classical complexity-zoo to include quantum classes. The existence of algorithms such as Shor's algorithm[4] motivates the addition of additional quantum classes to the complexity zoo. We exemplify with the two most common quantum complexity classes. Additionally, in Figure 4 we illustrate a commonly believed view of the complexity zoo.

Example 3.20 *Bounded-error quantum polynomial time (BQP).*
The first such class is bounded-error quantum polynomial time, and is the quantum analogy of bounded-error probabilistic polynomial time (BPP). BQP is the class of problems for which there exists a polynomial time randomised quantum algorithm that solves every instance with a success probability of at least 2/3. Shor's algorithm is contained in BQP.

Example 3.21 *Quantum Merlin-Arthur (QMA).*
The Quantum analogy of NP is quantum Merlin-Arthur. This is the set of problems, where if a specific instance of a problem is a yes-instance, then there is a polynomial size quantum proof (a quantum state), that the polynomial time verifier accepts with high probability. Additionally, when the answer is no, every polynomial size quantum "proof" is rejected with high probability.

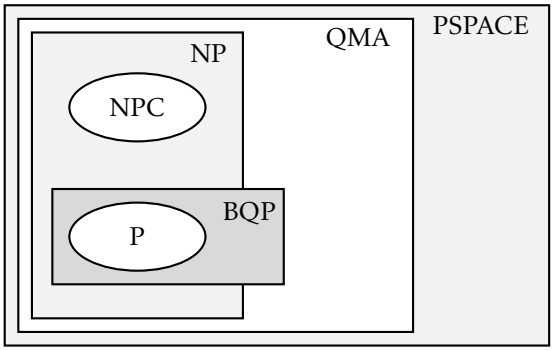

Figure 4: A view of the complexity zoo

So far, we have omitted discussing the model of computation that is the quantum analogy of the Turing Machine, and in our statement of the QECTT we just said "quantum computer". In practice, there is a direct analogy to the classical Turing Machine — the Quantum Turing Machine — but it is more common to use the quantum circuit model. This model consists of a sequence of quantum gates, acting sequentially on n qubits, i.e. $(\mathbb{C}^2)^{\otimes n}$. As gates one allow any unitary transformation on $(\mathbb{C}^2)^{\otimes n}$. Note that a gate that only acts a subset of the n qubits is extended to $(\mathbb{C}^2)^{\otimes n}$ by letting it be the identity transformation on the remaining qubits.

[4]Shor's algorithm factors an integer in polynomial time, a problem that is generally not believed to be solvable in polynomial time on a classical computer.

At this point, we can define when a quantum algorithm runs in polynomial-time. We say that a quantum algorithm is polynomial in n, the size of its input, if it is implementable by a family of quantum circuits $\{C_n\}_{n\in\mathbb{N}}$, such that the number of gates in C_n is polynomial in n. With this definition we can define when two quantum states are quantum-computationally indistinguishable.

Definition 3.22 *Quantum-computationally indistinguishable.*
Two quantum states of size n, ρ and σ, are said to be quantum-computationally indistinguishable, if any polynomial in n algorithm has negligible in n advantage for distinguishing ρ from σ. We write this as

$$\rho \overset{qc}{\approx} \sigma.$$

3.3 Quantum information theory

Quantum information theory is the quantum analogy of classical information theory, and includes the adaptation of many topics that are studied there. Hence, quantum information theory includes the quantum version of topics such as noise, distinguishability, error-correction, entropy, and compression. In this section we present some quantum information theory, focusing mainly on entropy and comparing the closeness of quantum states.

One of the most fundamental results for quantum information theory is the no-cloning theorem, which states that it is not possible to perfectly duplicate a quantum state. The importance of this result cannot be overstated, and we will make both explicit and implicit use of it many times.

Theorem 3.23 *No-cloning theorem.*
There exists no state $|s\rangle \in \mathcal{H}$ and unitary U acting on $\mathcal{H}^2$, such that for every state $|\phi\rangle$

$$U(|\phi\rangle \otimes |s\rangle) = |\phi\rangle \otimes |\phi\rangle,$$

Proof. We assume that the $|s\rangle$ is an arbitrary, but fixed standard pure state. Assume towards a contradiction that such U exists. Then, in particular, we have that for any $|\phi\rangle$ and $|\psi\rangle$ in $\mathcal{H}$

$$\begin{aligned} U(|\phi\rangle \otimes |s\rangle) &= |\phi\rangle \otimes |\phi\rangle, \\ U(|\psi\rangle \otimes |s\rangle) &= |\psi\rangle \otimes |\psi\rangle. \end{aligned}$$

Taking the inner product of the right-hand sides we obtain that

$$\langle |\phi\rangle \otimes |\phi\rangle, |\psi\rangle \otimes |\psi\rangle \rangle = \langle \phi|\psi\rangle\langle\phi|\psi\rangle = (\langle\phi|\psi\rangle)^2,$$

and from the left-hand sides that

$$\begin{aligned} \langle U(|\phi\rangle \otimes |s\rangle), U(|\psi\rangle \otimes |s\rangle)\rangle &= \langle |\phi\rangle \otimes |s\rangle, U^* U(|\psi\rangle \otimes |s\rangle)\rangle \\ &= \langle |\phi\rangle \otimes |s\rangle, |\psi\rangle \otimes |s\rangle \rangle \end{aligned}$$

$$= \langle \phi | \psi \rangle \langle s | s \rangle$$
$$= \langle \phi | \psi \rangle .$$

Since the only solutions to $x = x^2$ is 0 and 1, this implies that either $|\phi\rangle = |\psi\rangle$, or $|\phi\rangle \perp |\psi\rangle$. But clearly this cannot always be the case, contradicting that U clones arbitrary states. As a concrete example of this failing, consider $\mathcal{H} = \mathbb{C}^2$, $|\phi\rangle = |0\rangle$ and $|\psi\rangle = |+\rangle$. Then neither $|\phi\rangle = |\psi\rangle$ nor $|\phi\rangle \perp |\psi\rangle$ are true, giving us a contradiction. □

The next topic for consideration is that of comparing quantum states. The canonical way of measuring the closeness of two states is by their trace distance, which defines a metric on $\mathbb{B}(H)$. In general, it is not as straight forward to compare quantum states as it is to compare classical states. Consider that in a classical world, talking about the closeness of two bit-strings can be done in several well-defined, and meaningful ways (depending on application). One could look at the number of indices in which they differ or compare their respective Hamming weight. When moving to a quantum setting where we might have superpositions of strings it is not immediately clear how these notions should even translate. We will consider this problem in Section 3.4.

Definition 3.24 *Trace norm and trace distance.*
For any state ρ, the trace norm of ρ is defined as

$$\|\rho\|_{\mathrm{tr}} := \mathrm{tr}\sqrt{\rho\rho^*}.$$

For any two states, ρ and σ, we define their trace distance as

$$\Delta(\rho, \sigma) := \frac{1}{2}\|\rho - \sigma\|_{\mathrm{tr}}.$$

It can be shown that an equivalent definition of the trace distance is

$$\Delta(\rho, \sigma) = \max_P \mathrm{tr}\, P(\rho - \sigma),$$

where the max is taken over all positive operators $P \leq \mathbb{I}$. This interpretation is harder to work with, but particularly useful for *understanding* the results we will obtain; it shows that the trace distance is exactly the maximal likeliness of distinguishing between ρ and σ. Hence, it must be the case that if $\Delta(\rho, \sigma) < \varepsilon$, then ρ acts exactly like σ, except with probability ε. This motivates the following definition.

Definition 3.25 *Statistically indistinguishable.*
Two quantum states ρ and σ are called (statistically) indistinguishable if their trace distance is negligible.

We note that this is a stronger property than being quantum-computationally indistinguishable, since for any circuit to be able to distinguish the states, they need to act differently. When we are only interested in two states being quantum-

computational indistinguishable we might therefore just write $\rho \overset{qc}{\approx} \sigma$, even if they are in fact statistically indistinguishable.

For details about this interpretation of the trace distance, and a proof of the equivalence of the definitions, we refer to chapter 9.2.1 of [NC02] and 2.11 of [Gra21].

An useful identity for computing the trace distance of pure states is given by the following theorem.

Theorem 3.26 *An identity for the trace distance of pure states.*
If $|\phi\rangle$ and $|\psi\rangle$ are both pure states the following identity holds:

$$\Delta(|\phi\rangle\langle\phi|, |\psi\rangle\langle\psi|) = \sqrt{1 - |\langle\phi|\psi\rangle|^2}.$$

Proof. For readability we use the following abbreviations:

$$\begin{aligned}\pi_x &:= |x\rangle\langle x|,\\ D &:= \pi_\phi - \pi_\psi.\end{aligned}$$

Since D is self-adjoint, it follows from eigendecomposition that there must exist a unitary operator U such that for the diagonal operator Λ, with the eigenvalues of D on the diagonal, we have that $D = U\Lambda U^*$. Therefore, we have that

$$\begin{aligned}\|D\|_{\mathrm{tr}} &\overset{\mathrm{def}}{=} \mathrm{tr}\sqrt{DD^*}\\ &= \mathrm{tr}\sqrt{U\Lambda U^* U\Lambda U^*}\\ &= \mathrm{tr}\sqrt{\Lambda^2}\\ &= \sum_i |(\Lambda)_{ii}|. \end{aligned} \tag{3.2}$$

I.e., the trace-norm of D is the sum of the absolute value of the eigenvalues of D. If $\pi_\phi = \pi_\psi$ the theorem is trivially true, since then the trace distance is 0, and $\langle\phi|\psi\rangle = 1$. Hence, we assume that $\pi_\phi \neq \pi_\psi$. Then the rank of D is 2, and therefore we have two non-zero eigenvalues. Now, using the properties of the trace we obtain that

$$\begin{aligned}\mathrm{tr}(A) &= \mathrm{tr}(\pi_\phi) - \mathrm{tr}(\pi_\psi)\\ &= \mathrm{tr}(|\phi\rangle\langle\phi|) - \mathrm{tr}(|\psi\rangle\langle\psi|)\\ &= \mathrm{tr}(\langle\phi|\phi\rangle) - \mathrm{tr}(\langle\psi|\psi\rangle)\\ &= 0.\end{aligned}$$

Since the trace of a matrix is the sum of its eigenvalues, we must have a pair $\pm\lambda$ for some $\lambda \in \mathbb{R}$. Together with Equation (3.2) this implies that $\|D\|_{\mathrm{tr}} = 2\lambda$. To find λ we now observe that λ^2 is an eigenvector for D^2, and hence

$$\begin{aligned}2\lambda^2 &= \mathrm{tr}\, D^2\\ &= \mathrm{tr}\,(\pi_\phi - \pi_\phi)^2\end{aligned}$$

$$= \operatorname{tr} \pi_\phi^2 - 2 \operatorname{tr} (\pi_\phi \pi_\psi) + \operatorname{tr} \pi_\psi^2 \tag{3.3}$$

$$= 2 - 2|\langle \phi | \psi \rangle|^2, \tag{3.4}$$

where (3.3) follows from linearity of the trace, and (3.4) from the states being pure, and the cyclic property of the trace. Note that this implies that $\lambda = \sqrt{1 - |\langle \phi | \psi \rangle|^2}$.

Finally, combining this with our earlier observation that $\|D\|_{\mathrm{tr}} = 2\lambda$, we obtain that

$$\|D\|_{\mathrm{tr}} = 2\sqrt{1 - |\langle \phi | \psi \rangle|^2},$$

which, together with the definition of the trace distance, proves the theorem. □

3.3.1 Entropy and privacy amplification

In this section we extend the notion of the entropy in a random variable to the entropy in a quantum system. Our focus is on the min-entropy for classical-quantum systems, which we will make use of later, but we also introduce von Neumann entropy, which is the quantum equivalence to Shannon entropy, defined in Definition 2.3.

Similarly to how the Shannon entropy can be considered a measure of the randomness in a random variable, the von Neumann entropy can be considered a measure for how much uncertainty there is about the state of a quantum system.

Definition 3.27 *von Neumann entropy.*
For a density matrix $\rho \in \mathbb{B}(\mathcal{H}_A)$ with non-zero eigenvalues $\lambda_1, \ldots, \lambda_n$, the von Neumann entropy is defined as

$$S(\rho) := \sum_{i=1}^{n} \lambda_i \log \lambda_i.$$

Note that one would expect there to be no uncertainty about the state of a pure state, and indeed, a pure state $|\psi\rangle$ has only one non-zero eigenvalue, namely 1, and hence, the von Neumann entropy in $|\psi\rangle\langle\psi|$ is

$$S(|\psi\rangle\langle\psi|) = 1 \cdot \log 1 = 0.$$

While the notion of both Shannon and von Neumann entropy has many uses, it is of limited use for indicating security in cryptography. Consider the following example.

Example 3.28
Alice uses a probabilistic process to obtain a key $K \in \{0,1\}^n$, which she then uses to encrypt some data. However, the process have a 50% chance of returning the all-zero string, and a 50% chance of returning a non-zero string of length n, chosen uniformly at random. Considering the key as a random variable, the

Shannon entropy is easy to increase; for example, going from $n = 8$ to $n = 16$ we increase the Shannon entropy in K from $H(K) \simeq 5$ to $H(K) \simeq 9$. However, we note that there is still a 50% chance that the key is the all-zero string, so any attacker would be wise to try this key first, making the encryption essentially useless half the time. This clearly show that another notion of entropy could be useful.

For this purpose we introduce the notion of min-entropy. In a classical setting, the min-entropy is

$$\mathrm{H}_{\min}(P) = -\log(\max_{p_i \in P} p_i),$$

i.e. minus the logarithm of the most likely outcome. Hence, in Example 3.28 the min-entropy of K is exactly 1, regardless of the length of K. An equivalent definition is the largest $\alpha \in \mathbb{R}$ such that all events occur with probability at most $2^{-\alpha}$.

Moving to the quantum setting, one can define the min-entropy of a quantum system using the operator norm of the density matrix, that is

$$\mathrm{H}_{\min}(\rho) := -\log \|\rho\|_\infty,$$

where the sup-norm essentially takes the role of max in the classical setting. However, we will be interested in Renner's conditional min-entropy [Ren05] of a classical system X conditioned on a quantum system E.

Definition 3.29 *Conditional min-entropy.*
For a classical-quantum state ρ_{XE} ,the min-entropy of X conditioned on E is

$$\mathrm{H}_{\min}(\rho_{XE}|E) := \sup_{\sigma_E} \max\{h \in \mathbb{R} \colon 2^{-h}\mathbb{I}_X \otimes \sigma_E - \rho_{XE} \geq 0\}.$$

We note that just like the von Neumann entropy is similar to the Shannon entropy, the quantum min-entropy is similar to the classical min-entropy. We take supremum over $\sigma_E \in \mathcal{H}_E$ and then the largest $h \in \mathbb{R}$ such that $2^{-h}\mathbb{I}_X \otimes \sigma_E - \rho_{XE}$ is a positive semi-definite operator, which is conceptually the operator equivalent of a number being non-negative. Note that in the special case where σ_X is the maximally mixed state, one obtains that $h = \log d$ where d is the dimension of X. Additionally, we use the following fundamental property of the min-entropy; it can only increase when parts of the system it is condition on are measured. If σ'_E is obtained by measuring σ_E then

$$\mathrm{H}_{\min}(\rho_{XE}|\sigma_E) \leq \mathrm{H}_{\min}(\rho_{XE}|\sigma'_E). \tag{3.5}$$

An essential tool for when we prove the security of the OT-protocol, is privacy amplification. Conceptually, privacy amplification allow us to extract a short but highly secret string from a longer but only partially secret string. We use Renner's privacy amplification theorem [Ren05; RK05] for this, which relies on two-universal hash functions.

Definition 3.30 *Two-universal hash function.*
A family F of functions $f : \mathcal{X} \to \mathcal{Y}$ is called a two-universal hash function if for $f \leftarrow_R F$ and for any distinct $x, y \in \mathcal{X}$ the following holds:

$$\Pr[f(x) = f(y)] \leq \frac{1}{|\mathcal{Y}|}.$$

From now on, we will assume that $\mathcal{Y} = \{0,1\}^\ell$ for some $\ell \in \mathbb{N}$, and that $\mathcal{X} = \{0,1\}^n$ with $n > \ell$. Renner's theorem for privacy amplification then guarantees the following.

Theorem 3.31 *Universal hash functions and privacy amplification.*
Let ρ_{XE} be a classical-quantum state. Let F be a two-universal hash function from $\{0,1\}^n$ to $\{0,1\}^\ell$, and $f \leftarrow_R F$. For $K = f(X)$ we then have that

$$\Delta(\rho_{KE}, 2^{-\ell}\mathbb{I}_K \otimes \rho_E) \leq \frac{1}{2} \cdot 2^{-\frac{1}{2}\mathrm{H}_{\min}(X|E)-\ell}.$$

Which is in line with our description of what privacy amplification is; if X contains sufficiently much more than ℓ bits of entropy even when given E, then f extracts a nearly uniformly random string of length ℓ from X.

3.4 Sampling games in a quantum setting

In this section we will build upon the work we did in Section 2.3, and extend our notion of sampling strategies to the quantum setting, where we consider composite qubit systems, rather than bit strings. This section consists mostly of a presentation of the theory for quantum sampling as developed by Bouman and Fehr [BF10], with a greater amount of details and extended explanations of the proofs of Theorem 3.39 and Lemma 3.41.

Some of the work we did in Section 2.3 can be directly extended to a quantum setting, but other parts will need to be reconsidered. This is exemplified by the fact that while a sampling strategy can be defined exactly the same way,[5] it is not clear how to interpret the value obtained from this. While we want this to indicate if the remaining subsystem is close to an all-zero state, it is not clear what "close to" should mean, eg. should it be that only a few subsystems can be in states other than $|0\rangle$, or maybe that if the system is in a superposition of different states, then any state that is different from the all-zero state has a very small amplitude? Before we delve into answering this question (and more) we will describe our setting.

Remark 3.32
As noted in Footnote 5, a classical sampling strategy is also well-defined in the quantum setting, so Definition 2.16 is directly reused as our definition of

[5] Indeed, any classical sampling strategy can be directly adapted to a quantum sampling strategy; one simply measures the subsystems indexed by t in the computational basis to obtain a string $q_t \in \{0,1\}^{|t|}$.

a sampling strategy.

Where we previously considered bit strings of length n, we will now consider n-partite quantum systems $A = A_1 A_2 \cdots A_n$, where for simplicity we assume that each subsystem is a qubit, i.e. the state space of each subsystem is $\mathbb{C}^2$, and hence $\mathcal{H}_A = (\mathbb{C}^2)^{\otimes n}$. Further, we allow A to be entangled with some arbitrary state E, with finite dimensional state space $\mathcal{H}_E$. We assume that the joint state of AE is pure, and note that if this was note the case, one could purify the state by adding a system R as in Theorem 3.16, or equivalently by increasing the dimension of $\mathcal{H}_E$. Hence, we write the state of AE as $|\phi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$.

Turning our attention back to the problem of what property we want to be able to conclude about the unused systems in A, a natural approach is to first extend the definition of relative Hamming weight, and then use this definition to describe the desired property.

Definition 3.33 *Relative Hamming weight of a classical-quantum state.*
For a system AE, a state $|\phi_{AE}\rangle$ of AE has relative Hamming weight β within A if it is of the form $|\phi_{AE}\rangle = |b\rangle|\phi_E\rangle$ with $b \in \{0,1\}^n$ and $\omega(b) = \beta$. Similarly, for any subsystem $A' \subseteq A$ the relative Hamming weight of AE within A' is $\omega(b')$ where b' is the string of the entries in b, corresponding to the systems in A'.

With this definition a natural property — and a useful one, as we will later see — is the following. If the outcome of a sampling strategy is $f(t, q_t, s)$, then we wish to conclude that the state of the remaining subsystem $A_{\bar{t}}E$ is in a superposition of states with relative Hamming weight within $A_{\bar{t}}$ close to the outcome, up to a (small) error. To do this, we (ab)use the set $B^{\delta}_{t,s}$, of strings $b \in \{0,1\}^n$ such that $|\omega(b_{\bar{t}}) - f(t, b_t, s)| < \delta$, see (2.1), by extending it to the quantum setting, defining

$$\operatorname{span}(B^{\delta}_{t,s}) := \operatorname{span}\{|q\rangle : q \in B^{\delta}_{t,s}\} = \operatorname{span}\{|q\rangle : |\omega(q_{\bar{t}}) - f(t, q_t, s)| < \delta\} \subseteq \mathcal{H}_A.$$

I.e., the subspace of vectors that are linear combinations of "strings" with relative Hamming weight δ-close to $f(t, q_t, s)$ within $A_{\bar{t}}$. This definition implies that if a state $|\phi_{AE}\rangle$ is in $\operatorname{span}(B^{\delta}_{t,s}) \otimes \mathcal{H}_E$ for some t and s, and exactly this t and this s are chosen when sampling, we can be certain that regardless of what we find $f(t, q_t, s)$ to be, the state of $A_{\bar{t}}E$ is in a superposition of states with relative Hamming weight δ-close to $f(t, q_t, s)$ within $A_{\bar{t}}$.

Example 3.34
Let us consider what elements $\operatorname{span}(B^{\delta}_{t,s})$ consists of linear combinations of. For simplicity, we restrict our consideration to the sampling strategy from Example 2.19, i.e. random sampling without replacement.

Clearly, for any δ, s, and t, we would have $|00\cdots0\rangle$ and $|11\cdots1\rangle$ in $\operatorname{span}(B^{\delta}_{t,s})$; the estimate of the Hamming weight, $f(t, q_t, s)$, is respectively 0 and 1, which is also exactly the Hamming weight of the remaining subsystem.

Now, what other elements are in there depends on t and δ. Assume for

simplicity that n is even, and $|t| = n/2$ (i.e. that exactly half of the entries are used for estimating). Then we consider some cases

- If $\delta < 1/|t|$, there are no strings with a single entry differing from the rest; for such a string either $f(t, q_t, s) = 0$ (or 1) and $\omega(q_{\bar{t}}) = 1/|t|$ (or $1 - 1/|t|$), or vice versa, meaning that $|\omega(q_{\bar{t}}) - f(t, q_t, s)| = 1/|t| > \delta$. On the other hand, for any fixed t, $\mathrm{span}\,(B^{\delta}_{t,s})$ contains all the strings with one differing entry in q_t and one differing entry in $q_{\bar{t}}$; here either $\omega(q_{\bar{t}}) = 1/|t| = f(t, q_t, s)$ or $\omega(q_{\bar{t}}) = 1 - 1/|t| = f(t, q_t, s)$.

  Similar arguments hold for all strings with an odd number of entries different from the majority, and for strings with the same number of 0's and 1's inside q_t and $q_{\bar{t}}$.

- If $1/|t| \leq \delta < 2/|t|$ we have not only the strings that we argued was in the set if $\delta < 1/|t|$, but now also all the elements where there is one more differing entry in either q_t or $q_{\bar{t}}$, eg. for any t we have $|100\cdots0\rangle$, $|010\cdots0\rangle$, and so on, in $\mathrm{span}\,(B^{\delta}_{t,s})$.

One could continue for $2/|t| \leq \delta < 3/|t|$, and so on, and for different sizes of t, where it starts to matter if the differing entry is in q_t or $q_{\bar{t}}$.

Having defined the property that we want to conclude that the remaining subsystem has, we turn our attention to defining the error probability in our quantum setting. Essentially, the idea is to take the worst case distance between the real state, and the closest ideal state, for which the property holds. To formalise this we use the following construction. Consider the classical-quantum system $TSAE$, given by the classical independent random variables T and S, with distributions P_T and P_S, and the quantum systems A and E as in the start of this section. Since the choice of t and s is independent of the state of AE, the state of $TSAE$ can be written as

$$\rho_{TSAE} = \sum_{t,s} P_{TS}(t,s)\,|t,s\rangle\langle t,s| \otimes |\phi_{AE}\rangle\langle\phi_{AE}|.$$

We will call this the *real state* of TSAE, and we wish to compare it to an *ideal state* of $TSAE$, which will be of the form

$$\tilde{\rho}_{TSAE} = \sum_{t,s} P_{TS}(t,s)\,|t,s\rangle\langle t,s| \otimes |\tilde{\phi}^{t,s}_{AE}\rangle\langle\tilde{\phi}^{t,s}_{AE}|, \tag{3.6}$$

where for $\delta > 0$, we require that $|\tilde{\phi}^{t,s}_{AE}\rangle \in \mathrm{span}\,(B^{\delta}_{t,s})$ for all t and s, and where $|\tilde{\phi}^{t,s}_{AE}\rangle$ therefore implicitly not required to be independent of t and s. Note that this ideal state has exactly the property we required of it, namely that after measuring A_t and observing q_t, the state of $A_{\bar{t}}E$ is in a superposition of classical-quantum states with relative Hamming weight δ-close to $f(t, q_t, s)$ within $A_{\bar{t}}$. We can now define the error probability in our quantum setting, but before doing so we give a small example.

Example 3.35 *Continuation of Example 3.34.*
Consider a 4-qubit system A, without an entangled system, and continue using the sampling strategy from Example 2.19, with $|t| = n/2 = 2$ and $\delta < 1/|t| = 1/2$. Here we could imagine that AE (or just A, since there is no entangled system E) is in the pure state $|1111\rangle$. Writing t as a bit-string with a 1 indicating that we use the entry for testing, and a 0 indicating that we do not, we obtain that the real state is

$$\begin{aligned}\rho_{TSAE} =& \rho_{TA} = \sum_t P_T(t)|t\rangle\langle t| \otimes |1111\rangle\langle 1111| \\ =& \tfrac{1}{6}|1100\rangle\langle 1100| \otimes |1111\rangle\langle 1111| + \tfrac{1}{6}|1010\rangle\langle 1010| \otimes |1111\rangle\langle 1111| \\ &+ \tfrac{1}{6}|1001\rangle\langle 1001| \otimes |1111\rangle\langle 1111| + \tfrac{1}{6}|0110\rangle\langle 0110| \otimes |1111\rangle\langle 1111| \\ &+ \tfrac{1}{6}|0101\rangle\langle 0101| \otimes |1111\rangle\langle 1111| + \tfrac{1}{6}|0011\rangle\langle 0011| \otimes |1111\rangle\langle 1111|,\end{aligned}$$

while an ideal state could be

$$\begin{aligned}\tilde{\rho}_{TSAE} =& \tilde{\rho}_{TA} = \sum_t P_T(t)|t\rangle\langle t| \otimes |\tilde{\phi}_A^t\rangle\langle \tilde{\phi}_A^t| \\ =& \tfrac{1}{6}|1100\rangle\langle 1100| \otimes |1001\rangle\langle 1001| + \tfrac{1}{6}|1010\rangle\langle 1010| \otimes |1100\rangle\langle 1100| \\ &+ \tfrac{1}{6}|1001\rangle\langle 1001| \otimes |0011\rangle\langle 0011| + \tfrac{1}{6}|0110\rangle\langle 0110| \otimes |1111\rangle\langle 1111| \\ &+ \tfrac{1}{6}|0101\rangle\langle 0101| \otimes |0011\rangle\langle 0011| + \tfrac{1}{6}|0011\rangle\langle 0011| \otimes |0000\rangle\langle 0000|.\end{aligned}$$

Note that each $\tilde{\phi}_A^t$ is in $\mathrm{span}(B_t^\delta)$. To lower bound the distance between ρ_{TA} and the ideal states, one could now take the trace distance between ρ_{TA} and $\tilde{\rho}_{TA}$.

Definition 3.36 *Quantum error probability.*
The quantum error probability of a sampling strategy $\Psi = (P_T, P_S, f)$ is parameterised by $0 < \delta < 1$ and defined as

$$\varepsilon_{\text{quant}}^\delta(\Psi) := \max_{\mathcal{H}_E} \max_{|\phi_{AE}\rangle} \min_{\tilde{\rho}_{TSAE}} \Delta(\rho_{PTSA}, \tilde{\rho}_{TSAE}), \tag{3.7}$$

where the first max is over all finite-dimensional state spaces $\mathcal{H}_E$, the second max is over all state vectors $|\phi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$, and the min is over all ideal states $\tilde{\rho}_{TSAE}$, defined as in (3.6).

Example 3.37 *Continuation of Example 3.35.*
Comparing the definition of the error probability and the states described in Example 3.35, we observe that comparing the two states we described would lower bound the two outer maxes; both the choice of E being the empty system, and our choice of $|\phi_{AE}\rangle$ is allowed. However, we did not pick the closest ideal state; since $|\phi_{AE}\rangle \in \mathrm{span}(B_{t,s}^\delta)$ we could have chosen $\tilde{\rho}_{TA} = \rho_{TA}$, with which we just obtain the rather trivial lower bound 0, for $\varepsilon_{\text{quant}}^\delta$

Summing up, this definition guarantees that even in the worst case, the state of $A_{\bar{t}}E$ is at least $\varepsilon_{\text{quant}}^\delta$-close to some superposition of states with relative Hamming weight δ-close to $f(t, q_t, s)$, and (remembering that an alternative

definition of the trace distance is the maximal probability of telling the two states apart) it therefore behaves *exactly* like such a superposition, except with probability $\varepsilon^{\delta}_{\text{quant}}$.

Remark 3.38
Similarly to the classical setting, all of these result generalise to comparing to any fixed reference state, rather than just the all-zero state. For instance, if $|\phi^{\circ}_A\rangle$ is given by

$$|\phi^{\circ}_A\rangle = |\hat{x}_1\rangle_{\hat{\theta}_1} \otimes |\hat{x}_2\rangle_{\hat{\theta}_2} \otimes \cdots \otimes |\hat{x}_n\rangle_{\hat{\theta}_n},$$

for fixed reference bases $\hat{\theta} \in \{+,\times\}^n$ and reference string $\hat{x} \in \{0,1\}^n$, then one could replace measuring in the computational basis by measuring in the Hadamard basis whenever $\hat{\theta}_i = \times$, and consider the relative Hamming distance to $\hat{x}$, instead of the relative Hamming weight.

3.4.1 Comparing quantum and classical error

While it might appear to be rather hard to compute the quantum error probability, we shall see that there is a simple connection between the error probability of a sampling strategy in the classical setting and in the quantum setting. Hence, when using the strategies we considered in Section 2.3.2 in a quantum setting, we can easily tell how well they are guaranteed to perform.

Theorem 3.39
For any sampling strategy Ψ and any $\delta > 0$ we have:

$$\varepsilon^{\delta}_{\text{quant}}(\Psi) \leq \sqrt{\varepsilon^{\delta}_{\text{class}}(\Psi)}.$$

Example 3.40 *Continuation of Example 3.37.*
Once we have proven this theorem we can immediately find an upper bound for the error of the situation considered in Example 3.35 and Example 3.37. Combining the theorem and Example 2.19 we see that

$$\varepsilon^{\delta}_{\text{quant}} \leq \sqrt{\varepsilon^{\delta}_{\text{class}}} \leq \sqrt{2\exp\left(-\frac{1}{2}\delta^2 2\right)}.$$

The proof of Theorem 3.39 consists of an explicit construction of an ideal state $\tilde{\rho}_{TSAE}$ with the desired properties, for an arbitrary real state. This also takes care of showing that (3.7) is actually well-defined, a consideration that we have so far omitted.

The idea behind the proof is relatively straight forward. Fixing a real state, we decompose it by projecting it onto the space of ideal states and its orthogonal compliment. One can then interpret the amplitudes of each of these two components as the probability that the measurement outcome of the real state is in each of the orthogonal subspaces, which gives a relation to the classical error.

Finally, using the identity from Theorem 3.26 and Jensen's inequality leads to the desired conclusion.

Proof of Theorem 3.39. We will show that for any fixed $\delta > 0$, arbitrary $\mathcal{H}_E$, and arbitrary $|\phi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$, there exists an ideal state $\tilde{\rho}_{TSAE}$ with

$$\Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE}) \leq \sqrt{\varepsilon^{\delta}_{\text{class}}}.$$

Towards this end, we fix $\delta > 0$, and for each pair t, s we define $|\tilde{\phi}^{ts}_{AE}\rangle$ to be the re-normalised orthogonal projection of $|\phi_{AE}\rangle$ onto $\text{span}(B^{\delta}_{t,s}) \otimes \mathcal{H}_E$, and $|\tilde{\phi}^{ts\perp}_{AE}\rangle$ to be the re-normalised orthogonal projection of $|\phi_{AE}\rangle$ onto $\text{span}(B^{\delta}_{t,s})^{\perp} \otimes \mathcal{H}_E$. Thus, we can decompose $|\phi_{AE}\rangle$ as

$$|\phi_{AE}\rangle = \langle\tilde{\phi}^{ts}_{AE}|\phi_{AE}\rangle|\tilde{\phi}^{ts}_{AE}\rangle + \langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|\tilde{\phi}^{ts\perp}_{AE}\rangle.$$

Note that due to the orthogonality of $|\tilde{\phi}^{ts}_{AE}\rangle$ and $|\tilde{\phi}^{ts\perp}_{AE}\rangle$ this implies that

$$\begin{aligned}
1 =& \||\phi_{AE}\rangle\|^2 \\
=& \|\langle\tilde{\phi}^{ts}_{AE}|\phi_{AE}\rangle|\tilde{\phi}^{ts}_{AE}\rangle + \langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|\tilde{\phi}^{ts\perp}_{AE}\rangle\|^2 \\
=& |\langle\tilde{\phi}^{ts}_{AE}|\phi_{AE}\rangle|^2\||\tilde{\phi}^{ts}_{AE}\rangle\|^2 + |\langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|^2\||\tilde{\phi}^{ts\perp}_{AE}\rangle\|^2 \\
=& |\langle\tilde{\phi}^{ts}_{AE}|\phi_{AE}\rangle|^2 + |\langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|^2. \qquad (3.8)
\end{aligned}$$

Observe that $|\langle\tilde{\phi}^{ts}_{AE}|\phi_{AE}\rangle|^2$ is exactly the probability of observing $q \in B^{\delta}_{t,s}$ when measuring subsystem A of $|\phi_{AE}\rangle$ in the computational basis, i.e. $\Pr[Q \in B^{\delta}_{t,s}]$, where Q is the random variable for the measurement outcomes. Similarly, $|\langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|^2$ is $\Pr[Q \notin B^{\delta}_{t,s}]$. Hence, we have that

$$\begin{aligned}
\sum_{t,s} P_{TS}(t,s)|\langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|^2 =& \sum_{t,s} P_{TS}(t,s)\Pr[Q \notin B^{\delta}_{t,s}] \\
=& \Pr[Q \notin B^{\delta}_{T,S}] \\
=& \sum_{q} P_Q(q)\Pr[q \notin B^{\delta}_{T,S}] \\
\leq& (\max_{q} \Pr[q \notin B^{\delta}_{T,S}]) \sum_{q} P_Q(q) \\
=& \max_{q} \Pr[q \notin B^{\delta}_{T,S}] \\
=& \varepsilon^{\delta}_{\text{class}}. \qquad (3.9)
\end{aligned}$$

If we now construct $\tilde{\rho}_{TSAE}$ as in (3.6), using $|\tilde{\phi}^{ts}_{AE}\rangle$ constructed as above, we can find the trace distance between ρ_{TSAE} and $\tilde{\rho}_{TSAE}$.

$$\begin{aligned}
\Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE}) =& \Delta\Big(\sum_{t,s} P_{TS}(t,s)|t,s\rangle\langle t,s| \otimes |\phi_{AE}\rangle\langle\phi_{AE}|, \\
& \sum_{t,s} P_{TS}(t,s)|t,s\rangle\langle t,s| \otimes |\tilde{\phi}^{t,s}_{AE}\rangle\langle\tilde{\phi}^{t,s}_{AE}| \Big) \qquad (3.10)
\end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{2} \Big\| \sum_{t,s} P_{TS}(t,s) |t,s\rangle\langle t,s| \otimes |\phi_{AE}\rangle\langle\phi_{AE}| \\
&\qquad - \sum_{t,s} P_{TS}(t,s) |t,s\rangle\langle t,s| \otimes |\tilde{\phi}^{t,s}_{AE}\rangle\langle\tilde{\phi}^{t,s}_{AE}| \Big\|_{\mathrm{tr}} && (3.11) \\
&\leq \frac{1}{2} \sum_{t,s} P_{TS}(t,s) \Big\| |t,s\rangle\langle t,s| \otimes |\phi_{AE}\rangle\langle\phi_{AE}| \\
&\qquad\qquad - |t,s\rangle\langle t,s| \otimes |\tilde{\phi}^{t,s}_{AE}\rangle\langle\tilde{\phi}^{t,s}_{AE}| \Big\|_{\mathrm{tr}} && (3.12) \\
&= \sum_{t,s} P_{TS}(t,s) \Delta\Big(|t,s\rangle\langle t,s| \otimes |\phi_{AE}\rangle\langle\phi_{AE}|, \\
&\qquad\qquad |t,s\rangle\langle t,s| \otimes |\tilde{\phi}^{t,s}_{AE}\rangle\langle\tilde{\phi}^{t,s}_{AE}| \Big) && (3.13) \\
&\leq \sum_{t,s} P_{TS}(t,s) \Delta(|\phi_{AE}\rangle\langle\phi_{AE}|, |\tilde{\phi}^{ts}_{AE}\rangle\langle\tilde{\phi}^{ts}_{AE}|) && (3.14)
\end{aligned}
$$

where Equation (3.10) follows from inserting the definitions of ρ_{TSAE} and $\tilde{\rho}_{TSAE}$, Equation (3.11) from the definition of the trace distance, Equation (3.12) from the triangle inequality, Equation (3.13) from going back to the trace distance, and Equation (3.14) from tracing out the classical systems and noting that the trace distance is contractive under trace preserving operations, such as that partial trace [NC02]. With some tricks we can further simplify this.

$$
\begin{aligned}
(3.14) &= \sum_{t,s} P_{TS}(t,s) \sqrt{1 - |\langle\tilde{\phi}^{ts}_{AE}|\phi_{AE}\rangle|^2} && (3.15) \\
&= \sum_{t,s} P_{TS}(t,s) \sqrt{|\langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|^2} && (3.16) \\
&= \mathrm{E}_{TS}\left[\sqrt{|\langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|^2} \right] && (3.17) \\
&\leq \sqrt{\mathrm{E}_{TS}\left[|\langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|^2 \right]} && (3.18) \\
&= \sqrt{\sum_{t,s} P_{TS}(t,s) |\langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|^2} && (3.19) \\
&\leq \sqrt{\varepsilon^{\delta}_{\mathrm{class}}}, && (3.20)
\end{aligned}
$$

where Equation (3.15) follows from both $|\phi_{AE}\rangle$ and $|\tilde{\phi}^{ts}_{AE}\rangle$ being pure, so we can apply the identity of Theorem 3.26. Equation (3.16) follows from Equation (3.8) implying that $|\langle\tilde{\phi}^{ts\perp}_{AE}|\phi_{AE}\rangle|^2 = 1 - |\langle\tilde{\phi}^{ts}_{AE}|\phi_{AE}\rangle|^2$, and Equations (3.17) and (3.19) are just packing and unpacking the definition of the expected value, so we can apply the concave version of Jensen's inequality (Theorem 2.2) in Equation (3.18), which we can since taking the square root is concave on the interval $[0,1]$. And finally, Equation (3.20) follows from Equation (3.9). □

Finally, two closing remarks on the theorem. First, we note that it has been shown that for many sampling strategies, e.g. Example 2.19, this bound is

tight [BF10]. Additionally, it should be noted that it follows from the Hilbert projection theorem that the ideal state $\tilde{\rho}_{TSAE}$ as constructed in the proof is indeed the state that minimises the distance to ρ_{TSAE}.[6]

3.4.2 Bounding the min-entropy

As the final part of our investigation into sampling strategies, we will see how we can use that a state is (close to) a superposition of states with approximately a given relative Hamming weight, to bound the min-entropy of a state. Our result is based on the following lemma, which essentially states that measuring parts of a superposition of a small number of orthogonal states produces a similar amount of uncertainty as when measuring the mixture of these orthogonal states, where entanglement can play no role.

Lemma 3.41
For A and E, arbitrary quantum systems, $\{|i\rangle\}_{i\in\mathcal{I}}$ and $\{|w\rangle\}_{w\in\mathcal{W}}$ be orthonormal bases of $\mathcal{H}_A$, and let $|\phi_{AE}\rangle$ and ρ_{AE}^{mix} have the forms

$$|\phi_{AE}\rangle = \sum_{i\in\mathcal{J}} \alpha_i |i\rangle|\phi_E^i\rangle,$$
$$\rho_{AE}^{\text{mix}} = \sum_{i\in\mathcal{J}} |\alpha_i|^2 |i\rangle\langle i| \otimes |\phi_E^i\rangle\langle\phi_E^i|,$$

for some $\mathcal{J} \subseteq \mathcal{I}$. If one let ρ_{WE} and ρ_{WE}^{mix} denote the classical-quantum systems obtained by measuring subsystem A of $|\phi_{AE}\rangle$ and ρ_{AE}^{mix}, respectively, in basis $\{|w\rangle\}_{w\in\mathcal{W}}$, observing outcome W, then

$$\text{H}_{\min}(\rho_{WE}|E) \geq \text{H}_{\min}(\rho_{WE}^{\text{mix}}|E) - \log|J|.$$

It is essentially sufficient to show that $\text{H}_{\min}(\rho_{WE}^{\text{mix}}|E) - \log|J|$ is contained in the set that $\text{H}_{\min}(\rho_{WE}|E)$ is the supremum of. However, showing this takes some work. What we do is that we first argue that if $|J|\rho_{WE}^{\text{mix}} - \rho_{WE}$ is positive semi-definite operator, then the claim holds. Then we show that this is indeed the case, by straight forwards algebraic calculations.

Proof. By Definition 3.29 we have to show that for $h = \text{H}_{\min}(\rho_{WE}^{\text{mix}}|E)$ and arbitrary $\sigma_E \in \mathcal{H}_E$ the following operator is positive semi-definite

$$2^{-(h-\log|J|)}\mathbb{I}_W \otimes \sigma_E - \rho_{WE}, \tag{3.21}$$

since then $\text{H}_{\min}(\rho_{WE}^{\text{mix}}|E) - \log|J|$ is in the set that we take the supremum over.

To show this, we will show that $|J|\rho_{WE}^{\text{mix}} \geq \rho_{WE}$ (in the sense that $|J|\rho_{WE}^{\text{mix}} - \rho_{WE}$, is positive semi-definite), since it then follows that

$$2^{-(h-\log|J|)}\mathbb{I}_W \otimes \sigma_E - \rho_{WE} = 2^{-h}|J|\mathbb{I}_W \otimes \sigma_E - \rho_{WE}$$

[6] The Hilbert projection theorem can for example be found as part of Theorem 1.24 of MacCluer's Elementary Functional Analysis [Mac09]. The claim follows immediately from point (d).

$$\begin{aligned}&\geq 2^{-h}|J|\mathbb{I}_W \otimes \sigma_E - |J|\rho_{WE}^{\text{mix}} \\ &= |J|(2^{-h}\mathbb{I}_W \otimes \sigma_E - \rho_{WE}^{\text{mix}}),\end{aligned} \tag{3.22}$$

and that the last line is positive semi-definite follows immediately from our assumption; h is picked such that $2^{-h}\mathbb{I}_X \otimes \sigma_E - \rho_{XE}^{\text{mix}}$ is positive semi-definite, and hence $|J|(2^{-h}\mathbb{I}_X \otimes \sigma_E - \rho_{XE}^{\text{mix}})$ is also positive semi-definite. Equation (3.22) being positive semi-definite implies that (3.21) must also be positive semi-definite.

To show that $|J|\rho_{WE}^{\text{mix}} \geq \rho_{WE}$ we first write out the post-measurement states of ρ_{WE} and ρ_{WE}^{mix}

$$\begin{aligned}\rho_{WE} &= \sum_{w\in\mathcal{W}} (|w\rangle\langle w| \otimes \mathbb{I}_E)\rho_{AE}(|w\rangle\langle w| \otimes \mathbb{I}_E) \\ &= \sum_{w\in\mathcal{W}} (|w\rangle\langle w| \otimes \mathbb{I}_E)|\phi_{AE}\rangle\langle\phi_{AE}|(|w\rangle\langle w| \otimes \mathbb{I}_E) \\ &= \sum_{w\in\mathcal{W}} (|w\rangle\langle w| \otimes \mathbb{I}_E)\left(\sum_{i\in\mathcal{J}} \alpha_i |i\rangle|\phi_E^i\rangle\right)\left(\sum_{j\in\mathcal{J}} \overline{\alpha_j}\langle\phi_E^j|\langle j|\right)(|w\rangle\langle w| \otimes \mathbb{I}_E) \\ &= \sum_{w\in\mathcal{W}}\sum_{i,j\in J} \alpha_i\overline{\alpha_j}|w\rangle\langle w|i\rangle\langle j|w\rangle\langle w| \otimes |\phi_E^i\rangle\langle\phi_E^j|, \\ \rho_{WE}^{\text{mix}} &= \sum_{w\in\mathcal{W}} (|w\rangle\langle w| \otimes \mathbb{I}_E)\rho_{AE}^{\text{mix}}(|w\rangle\langle w| \otimes \mathbb{I}_E) \\ &= \sum_{w\in\mathcal{W}} (|w\rangle\langle w| \otimes \mathbb{I}_E)\left(\sum_{i\in\mathcal{J}} |\alpha_i|^2 |i\rangle\langle i| \otimes |\phi_E^i\rangle\langle\phi_E^i|\right)(|w\rangle\langle w| \otimes \mathbb{I}_E) \\ &= \sum_{i\in J} |\alpha_i|^2 \sum_{w\in\mathcal{W}} |\langle w|i\rangle|^2 |w\rangle\langle w| \otimes |\phi_E^i\rangle\langle\phi_E^i|.\end{aligned}$$

Now, recalling the definition of being positive semi-definite, we need to show that $\langle\xi|(|J|\rho_{WE}^{\text{mix}} - \rho_{WE})|\xi\rangle \geq 0$ for all $|\xi\rangle \in \mathcal{H}_{\mathcal{W}} \otimes \mathcal{H}_E$. In order to do this we first consider $|\xi\rangle$ on the special form $|\xi\rangle = |v\rangle|\psi_E\rangle$ with $v \in \mathcal{W}$. Noting that $\langle v|w\rangle = 0$ for $v, w \in \mathcal{W}$ and $v \neq w$, we obtain that

$$\begin{aligned}\langle\xi|\rho_{WE}|\xi\rangle &= \sum_{i,j\in J} \alpha_i\overline{\alpha_j}\langle v|i\rangle\langle j|v\rangle\langle\psi_E|\phi_E^i\rangle\langle\phi_E^j|\psi_E\rangle \\ &= \left(\sum_{i\in J} \alpha_i\langle v|i\rangle\langle\psi_E|\phi_E^i\rangle\right)\left(\sum_{j\in J} \overline{\alpha_j}\langle j|v\rangle\langle\phi_E^j|\psi_E\rangle\right) \\ &= \left(\sum_{i\in J} \alpha_i\langle v|i\rangle\langle\psi_E|\phi_E^i\rangle\right)\overline{\left(\sum_{i\in J} \alpha_i\langle v|i\rangle\langle\psi_E|\phi_E^i\rangle\right)} \\ &= \left|\sum_{i\in J} \alpha_i\langle v|i\rangle\langle\psi_E|\phi_E^i\rangle\right|^2, \\ \langle\xi|\rho_{WE}^{\text{mix}}|\xi\rangle &= \sum_{j\in J} |\alpha_j|^2|\langle v|j\rangle|^2|\langle\psi_E|\phi_E^j\rangle|^2\end{aligned}$$

$$\geq \frac{1}{|J|}\left|\sum_{j\in J}\alpha_j \langle v|j\rangle\langle\psi_E|\phi_E^j\rangle\right|^2 \tag{3.23}$$
$$= \frac{1}{|J|}\langle\xi|\rho_{WE}|\xi\rangle,$$

where the inequality in (3.23) is an application of Titu's Lemma (Corollary 2.1). At this point the non-negativity in this special case follows from linearity

$$\begin{aligned}\langle\xi|(|J|\rho_{WE}^{\text{mix}} - \rho_{WE})|\xi\rangle &= \langle\xi||J|\rho_{WE}^{\text{mix}}|\xi\rangle - \langle\xi|\rho_{WE}|\xi\rangle \\ &\geq \langle\xi|\rho_{WE}|\xi\rangle - \langle\xi|\rho_{WE}|\xi\rangle = 0.\end{aligned}$$

For the general case we note that we can write any elements $|\xi\rangle$ as

$$|\xi\rangle = \sum_{w\in\mathcal{W}}\beta_w|w\rangle|\psi_E^w\rangle \in \mathcal{H}_{\mathcal{W}}\otimes\mathcal{H}_E,$$

and the claim now follows from linearity, and noting that for $v \neq v'$ we have

$$\langle v,\psi_E|\rho_{WE}|v',\psi_E'\rangle = 0 = \langle v,\psi_E|\rho_{WE}^{\text{mix}}|v',\psi_E'\rangle,$$

so all the cross-products vanish. □

We will now consider how we can apply this to an n-qubit system A that *is exactly* like we ensure that $A_{\bar{t}}$ *acts* (except with some small error) after we have measured A_t, i.e. where $|\phi_{AE}\rangle$ is a superposition of states with relative Hamming wight δ-close to β within A;

$$|\phi_{AE}\rangle = \sum_{\substack{b\in\{0,1\}^n \\ |\omega(b)-\beta|\leq\delta}}\alpha_b|b\rangle|\phi_E^b\rangle.$$

First note that for $\delta+\beta \leq \frac{1}{2}$ we have

$$|\{b\in\{0,1\}^n\colon |\omega(b)-\beta|\leq\delta\}| \leq |\{b\in\{0,1\}^n\colon \omega(b)\leq\delta+\beta\}| \leq 2^{\mathrm{h}(\beta+\delta)n},$$

where h is the binary entropy function, defined in Definition 2.3. Now, foreseeing that we have $J = \{b\in\{0,1\}^n\colon |\omega(b)-\beta|\leq\delta\}$, we see that we will have that $\log|J| \leq \mathrm{h}(\beta+\delta)n$.

Further, we note that since measuring the qubits in b in the Hadamard basis produces uniformly random bits, we can lower bound the min-entropy of the mixed state; from its construction there is no entanglement between A and E, and hence someone in control of E cannot know anything about the outcome of measuring these qubits, other than it was uniformly random. Assume that we measure $m \leq n$ of the bits in b in the Hadamard basis. Then there are at least m bits that are uniformly random, so even the most likely event cannot happen with probability greater than 2^{-m}, and hence $\mathrm{H}_{\min}(\rho_{AE}^{\text{mix}}|E) \geq m$.

The following corollary follows immediately from these observations and Lemma 3.41.

Corollary 3.42
Let A be an n-qubit system, let the state $|\phi_{AE}\rangle$ of AE be a superposition of states with relative Hamming weight δ-close to β within A, where $\delta + \beta \leq \frac{1}{2}$, and let X be the random variable obtained by measuring A in basis $\theta \in \{+, \times\}^n$, where m of the entries in θ are $\times$. Then

$$\mathrm{H}_{\min}(\rho_{XE}|E) \geq m - h(\beta + \delta)n.$$

Observe that this corollary is directly applicable to quantum cryptography, in the following way. Suppose that Bob prepares and hands over an n-qubit quantum system A to Alice. This system should be in the all-zero state (or any other state, see Remark 3.38), but since Alice cannot know if Bob is dishonest, the state could be anything, possibly even entangled to some system E that Bob controls. From Corollary 3.42 it then follows that if Alice applies a sampling strategy to convince herself that the remaining subsystem of AE is close to a superposition of states with bounded relative Hamming weight, then the measurement outcome of the remaining subsystem of A behaves similarly to the case where Bob honestly prepared A — and most importantly did not entangle A and E — with respect to the min-entropy. In Section 5.2 we will use this approach to show that the BBCS92 protocol is secure against a corrupted receiver.

4 Introduction to quantum cryptography

In this section we will introduce the notion of security that we use, and consider some introductory examples for quantum cryptography.

Before introducing our notion of security, we remark that it is not the only notion of security in a quantum world, and that there are even many problems where this definition is meaningless, guaranteeing nothing useful. This is due to the fact that we are only concerned with security against one of the participating parties acting dishonestly, and not in other forms of security, such as security against eavesdroppers. Hence, even for the classical quantum cryptographic result quantum key distribution, our notion of security is useless, and a different notion should be used.

It is a well known result that when we consider adversaries that can perform quantum computations, many well known, and widely used, classical computational hardness assumptions are no longer secure. Probably the best know example of this is Shor's algorithms that factors integers and compute discrete logarithms in polynomial time, implying that many classical computational hardness assumptions no longer hold (eg. the RSA and discrete logarithm assumptions). From here there are two different paths to consider. One is to replace the use of these broken assumptions with assumptions that are still believed to hold against quantum adversaries leading to what is referred to as "post-quantum cryptography". An example of such an assumption is the following.

Remark 4.1 *Learning With Errors (LWE).*
Let $\mathbb{Z}_q$ denote the ring of integers modulo q, and $\mathbb{Z}_q^n$ the set of vectors of size n over $\mathbb{Z}_q$. Further, let $f : \mathbb{Z}_q^n \to \mathbb{Z}_q$ be a linear functional. Then the LWE problem is to find f (or a close approximation) given a sample consisting of pairs (x, y) where $x \in \mathbb{Z}_q^n$ and $y \in \mathbb{Z}_q$ with $y := f(z) + \eta$ where the noise term η is sampled from some known distribution over $\mathbb{Z}_q$.

This problem was introduced by Oded Regev in 2005 [Reg05], who showed that the problem is as hard to solve as several worst-case lattice problems and how the problem could be used to create a public-key encryption scheme. The assumption that the LWE problem is hard to solve is referred to as the LWE-assumption, and is believed to hold against quantum adversaries.

A different approach is to explore how quantum mechanics can be used in a way that is beneficial to us. Typically, this is done by proving results under quantum versions of much weaker assumptions, when compared to the classical assumptions that the results usually rely on. This approach is referred to as quantum cryptography, and is where our focus is. We consider examples in this in Section 4.3, where we argue that there is an information-theoretically secure protocol for key distribution if one of the parties can send qubits to the other. Similar key-distribution protocols in a purely classical setting typically

relies on computational hardness assumptions. In Section 5 we show how to obtain oblivious transfer from bit-commitments, which in turn can be obtained from post-quantum one-way functions (pq-OWF), which is a much weaker assumption than the RSA-assumption that is typically used for oblivious transfer in a classical setting.

Definition 4.2 *Post-quantum one-way functions.*
Post-quantum one-way functions are one-way functions as defined in Definition 2.7, with the additional property that they are still secure against adversaries with quantum powers.

As indicated above, the assumptions used in quantum cryptography varies. Ideally, one would like to make as few assumptions as possible and obtain unconditionally secure protocols, like the one for quantum key-distribution. However, it has been shown that it is not possible to obtain unconditionally secure bit-commitments and oblivious transfer [Lo97]. However, both bit-commitments and oblivious transfer can be constructed from pq-OWFs, which can be considered the second weakest assumption. Drawing inspiration from Russell Impagliazzo's five worlds [Imp95], it is common to refer to the setting where pq-OWFs are assumed to exists as MiniQCrypt.

4.1 The notion of security

We use an extension of the classical ideal-/real-world paradigm, that is typically used in secure multiparty computation and was introduced in Section 2.2.1. Since the protocols we will consider are two-party protocols, we continue to state our security definitions in a two-party setting, but we note that they generalise to any (finite) number of players.

Our definition is essentially the same as Definition 2.9, except that Π now consists of two interactive polynomial time *quantum* machines (with classical output), and similarly the adversary is an interactive polynomial time *quantum* machines (with quantum output). Hence, our simulator must also have quantum output. Note that $\mathcal{F}$ is still classic; we do not consider quantum functionalities.

We also reuse the notion of the hybrid model, but stress that the composition theorem we use (Theorem 2.11), we only allow sequentially composition, i.e. only one subroutine can be active at a time. For a proof of the composition theorem in the quantum setting we refer to [HSS15].

Definition 4.3 *Quantum computational security against corrupted A.*
A protocol $\Pi = (A, B)$ implements a classical functionality $\mathcal{F}$ with quantum-computational security against corrupted A, if for any real-world (quantum) adversary $\mathcal{A}$, there exists an ideal-world simulator Sim, such that for any input to A and B it holds that the joint output in the real- and ideal-world

are quantum-computationally indistinguishable, i.e.

$$\mathrm{Out}_{\Pi,\mathcal{A},B} \overset{qc}{\approx} \mathrm{Out}_{\mathcal{F},\mathsf{Sim},\hat{B}}.$$

Equivalently, we say that there is no quantum distinguisher D that can distinguish between the situation pictured in Figure 5, and the situation pictured in Figure 6, with non-negligible advantage.

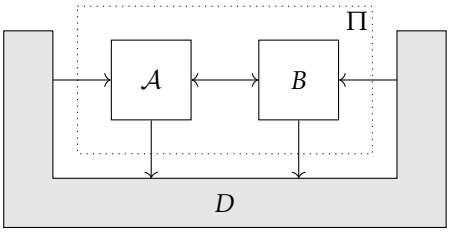

Figure 5: Illustration of the interactions in the real-world.

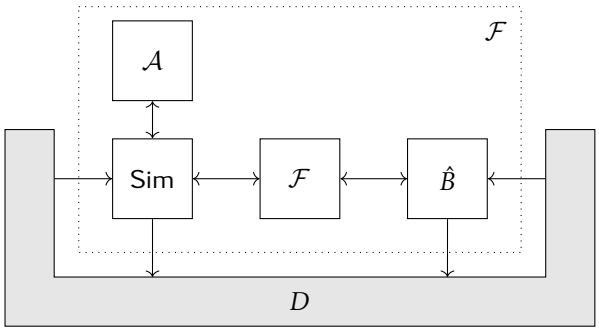

Figure 6: Illustration of the interactions in the ideal-world. Note that we have included an interaction between Sim and $\mathcal{A}$, to illustrate that Sim is allowed to use $\mathcal{A}$ as a subroutine.

We say that Π implements $\mathcal{F}$ with quantum-computational security, if it quantum-computationally securely implements $\mathcal{F}$ against both corrupted A and corrupted B. When case where the joint output of $\mathcal{A}$ and B is statistically indistinguishable from the joint output of Sim and B we say that Π implements $\mathcal{F}$ with unconditional security.

4.2 Quantum one-time-pad

As our first concrete example of quantum cryptography, we show how the one-time-pad can be extended to a quantum version, hiding an arbitrary pure qubit with density operator ρ, using two (classical) key bits, $k_1, k_2 \in \{0,1\}$. The key observation for this is that the Pauli matrices are their own inverses, so eg.

$XZZX = \mathbb{I}$. Hence, we can encrypt a qubit as follows

$$\mathsf{Enc}_{k_1,k_2}(\rho) = X^{k_1} Z^{k_2} \rho Z^{k_2} X^{k_1} =: \sigma,$$

and decrypt it

$$\mathsf{Dec}_{k_1,k_2}(\sigma) = Z^{k_2} X^{k_1} \sigma X^{k_1} Z^{k_2} = \rho.$$

This is clearly a case where the notion of security from Section 4.1 is not applicable, but we will show that if one does not know the keys, then σ reveals nothing about ρ, which intuitively show that the encryption is (perfectly) hiding.

Towards this we consider the joint classical-quantum state of the output of the encryption, and what keys are chosen. Denote by P the classical system and by A the quantum system. Then

$$\begin{aligned}\sigma_{PA} &= \sum_{i,j\in\{0,1\}} p_{ij} |ij\rangle\langle ij| \otimes \mathsf{Enc}_{i,j}(\rho) \\ &= \sum_{i,j\in\{0,1\}} \frac{1}{4} |ij\rangle\langle ij| \otimes X^i Z^j \rho Z^j X^i,\end{aligned}$$

where we have assumed that the keys are chosen uniformly at random. Now, to see how this state appears to an adversary, without knowledge of what keys have been chosen, we can trace out the classical system:

$$\begin{aligned}\sigma_A &:= \mathrm{tr}_P \, \sigma_{PA} \\ &= \frac{1}{4} \sum_{i,j\in\{0,1\}} \mathrm{tr}\left(|ij\rangle\langle ij|\right) X^i Z^j \rho Z^j X^i \\ &= \frac{1}{4} \sum_{i,j\in\{0,1\}} X^i Z^j \rho Z^j X^i \\ &= \frac{1}{4} \left(\rho + Z\rho Z + XZ\rho ZX + X\rho X\right).\end{aligned}$$

Finally, letting $\rho = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix}$, and observing that since ρ is pure, it must be the case that $\mathrm{tr}\,\rho = 1 = \alpha + \delta$, we have that

$$\begin{aligned}\sigma_A &= \frac{1}{4} \left(\rho + Z\rho Z + XZ\rho ZX + X\rho X\right) \\ &= \frac{1}{4} \left(\begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} + \begin{pmatrix} \alpha & -\beta \\ -\gamma & \delta \end{pmatrix} + \begin{pmatrix} \delta & -\gamma \\ -\beta & \alpha \end{pmatrix} + \begin{pmatrix} \delta & \gamma \\ \beta & \alpha \end{pmatrix} \right) \\ &= \frac{1}{4} \begin{pmatrix} 2(\alpha+\delta) & 0 \\ 0 & 2(\alpha+\delta) \end{pmatrix} \\ &= \frac{1}{2} \mathbb{I},\end{aligned}$$

i.e. regardless of the state of the qubit, the state of the "cipher text" appears to be maximally mixed, and thus an adversary cannot learn anything about the

input from it.

4.3 Quantum key-distribution

One of the most remarkable results in quantum cryptography is the unconditionally secure key-distribution protocol by Bennett and Brassard [BB84]. This protocol uses the unclonability of quantum states (see Theorem 3.23) to achieve information theoretic security, for a task that classically has to rely on computational assumptions. We will here introduce the protocol, and give some intuition for why it is secure, but not a full proof. An interesting note, however, is that one way to formally prove that it is secure is to use the sampling theory developed in Section 3.4 [BF10].

Definition 4.4 *BB84 states.*
A key part of the protocol is the use of qubits that are prepared in one of four states, with equal probability. These qubits are prepared as either $|0\rangle, |1\rangle, |+\rangle$ or $|-\rangle$, and we refer to qubits that are prepared this way as BB84 states. We note that one way to generate a BB84 state is to draw two classical bits, b_1 and b_2, uniformly at random, and then compute

$$|\phi\rangle = H^{b_1} X^{b_2} |0\rangle.$$

We will generally want many BB84 states, and when preparing them we will simply say that one draws random bits $x \leftarrow_R \{0,1\}^n$ and bases $\theta \leftarrow_R \{+,\times\}^n$ and then prepares n BB84 states $|x\rangle_\theta$, to be understood as

$$|x\rangle_\theta = |x^1\rangle \otimes |x^2\rangle \otimes \cdots \otimes |x^n\rangle,$$

where $|x^i\rangle$ is defined as

$$|x^i\rangle := H^{\theta_i} X^{x_i} |0\rangle,$$

interpreting $+$ as a 0 and $\times$ as a 1.

At this point, we can describe the first version of the protocol. The essential idea is that one part, Alice, prepares n BB84 states, which she then sends to the other party, Bob. Bob measures the qubits with respect to his own random choice of bases. Bob and Alice now exchange their respective choice of bases over a classical channel, and use the bits corresponding to the entries that they measured in the same bases as their key. We illustrate this in Figure 7. We assume that the classical channel is authenticated but not encrypted[7] and that the quantum channel is under complete control of the adversary.

Correctness of the protocol follows immediately from the fact that if Bob measures a qubit with respect to the same basis that Alice prepared it in, then Bob obtains the same bit that Alice used in the preparation.

[7] Due to this assumption this protocol is sometimes referred to as a key-expansion protocol, rather than a key-distribution.

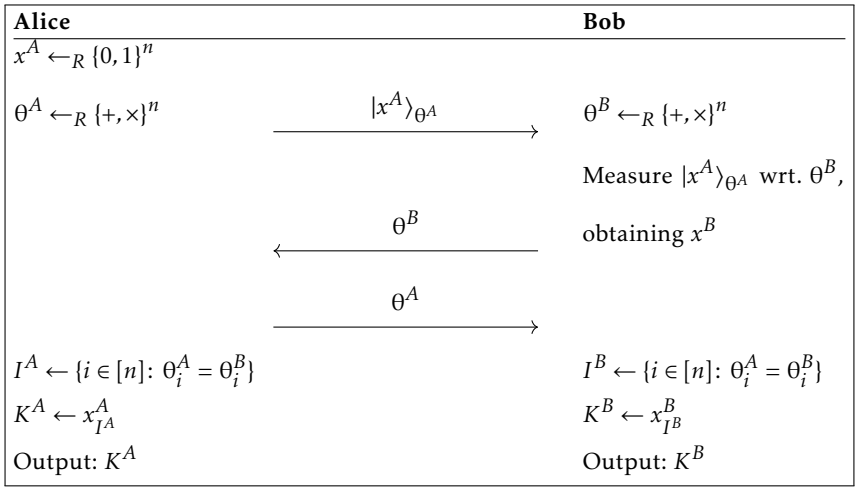

Figure 7: The concept behind quantum key-distribution.

While this protocol is not secure (we outline why below), we can already argue how the unclonability of quantum states is used. If the adversary can only eavesdrop on the channels, the unclonability theorem implies that the only action she can take is to measure some of the BB84 states as they are being transmitted. If she measures in any other basis than the one Alice prepared the state in, there is some non-zero probability that she obtains the wrong bit, when compared to what Alice prepared. Over all the roughly $n/2$ BB84 states that Alice and Bob measured in the same bases (and hence output as keys) it is therefore very unlikely that the adversary obtained all the same bits as the ones Alice used to prepare the $n/2$ BB84 states.

While the above argument is sound, it still leaves numerous holes in the security and correctness of the protocol, three of which we outline below.

1. First and foremost, it should be noted that if the adversary measures a qubit with respect to a different basis than the one Alice prepared the qubit in, the post-measurement state of the qubit is not the same as what Alice prepared. Hence, Bob's measurements might be different from Alice's, even when they measure with respect to the same basis. Hence, it need not be the case that $K^A = K^B$, and thus we can no longer be sure that the protocol enjoys correctness.

2. Once Alice broadcasts what bases she used to prepare the BB84 states the adversary might see that for some of the qubits, the bases it measured with respect to agree with what Alice prepared the qubits in. If Bob's basis also agrees with Alice's, this is a bit they will use for their respective keys, and hence the adversary knows some of the key.

3. Finally, if the adversary is not limited to eavesdropping, one could imagine that it stores all the BB84 states, and create some new states that it then

sends to Bob. When Alice then broadcast her bases, it measures the BB84 states in the same bases, learning the entirety of Alice's key.

To avoid the second problem, we use a privacy amplification function f, in the sense of Theorem 3.31, which we will not go further into details about here. For the third point, we introduce a test for errors, where Alice and Bob communicates classically to reveal if there is a difference between the BB84 states that Alice prepared, and the states that Bob received.

The test works as follows. Alice and Bob publish and compare the values that they have chosen/measured for roughly half the bits their bases agree on. We do this by using all the bits that they prepared/measured in the computational basis. Intuitively, this makes the protocol secure, by forcing the adversary to either measure in the computational basis, so as to not disturb the states they compare the outcome of (in which case she learns nothing about the bits they use for the key), or alternatively, learn the value of (some of) the bits they use for their shared key, but have a high chance of disturbing the states they compare, leading to Alice and Bob aborting with high probability. In Figure 8 we illustrate a protocol that implements these ideas.

We note that while this protocol is intuitively secure against an adversary learning the key outputted by either Alice or Bob, it is still vulnerable to the problem outlined in the first point. If the adversary measures every state in the computational basis the test will not detect anything, but there will be no correlation between Alice and Bob's keys. One way to fix this would be to instead have Alice choose a random subset of size $n/4$ of the states where their bases agreed, that they then run the verification on. In this case, the intuition for why the protocol is secure is that the odds of an attacker choosing the correct basis for enough of the $n/4$ qubits is negligible, and hence any tempering of the qubits will be detected with very high probability.

Alice		**Bob**
$x^A \leftarrow_R \{0,1\}^n$		
$\theta^A \leftarrow_R \{+,\times\}^n$	$\xrightarrow{\lvert x^A\rangle_{\theta^A}}$	$\theta^B \leftarrow_R \{+,\times\}^n$
		Measure $\lvert x^A\rangle_{\theta^A}$ wrt. θ^B,
	$\xleftarrow{\theta^B}$	obtaining x^B
	$\xrightarrow{\theta^A}$	
$J^A \leftarrow \{j \in [n] \colon \theta^A_j = \theta^B_j = +\}$		$J^B \leftarrow \{j \in [n] \colon \theta^A_j = \theta^B_j = +\}$
	$\xleftarrow{x^B_{J^B}}$	
	$\xrightarrow{x^A_{J^A}}$	
Verify that $x^B_{J^B} = x^A_{J^A}$		Verify that $x^B_{J^B} = x^A_{J^A}$
$I^A \leftarrow \{i \in [n] \colon \theta^A_i = \theta^B_i = \times\}$		$I^B \leftarrow \{i \in [n] \colon \theta^A_i = \theta^B_i = \times\}$
$K^A \leftarrow f(x^A_{I^A})$		$K^B \leftarrow f(x^B_{I^B})$
Output: K^A		Output: K^B

Figure 8: The secure version of quantum key-distribution.

5 Oblivious transfer in a quantum world

The archetypal protocol for oblivious transfer in a quantum world was presented by Bennet, Brassard, Crépeau, and Skubiszewska at CRYPTO 1991, and it will be refereed to as BBCS92 [Ben+91]. In this protocol, two parties (the Sender S and the Receiver R) first use quantum communication for the Sender to send a number of BB84 states to the Receiver, who then measures them in one of two bases, chosen at random. After this, they use post-quantum cryptography, classical communication, and classical computation, with values obtained from the measurements on the BB84 states, to finish up the OT.

It took nearly 20 years for the protocol to be proven secure against arbitrary quantum adversaries in the typical ideal-/real-world paradigm, and was finally done by Damgård et al. in 2009 [Dam+09]. However, in their proof they assume a commitment scheme with some special properties, which at that time was only known to be obtainable under strong assumptions, such as the LWE-assumption. The most recent development is that it has been shown by Grilo et al. [Gri+20] that we can construct commitments with the needed properties, using only post-quantum one-way functions, which implies that OT is in MiniQCrypt.

In Section 5.1 we present the BBCS92 protocol, and in Section 5.2 we show that it is secure, assuming ideal commitments. In Section 5.3 we consider the task of constructing commitments with the needed properties.

5.1 The protocol

The version of BBCS92 that was originally presented by Bennett et al. [Ben+91] is clearly insecure if the adversary is able to store quantum states for any meaningful amount of time. To remove this constraint, the authors suggest using commitments and cut-and-choose to verify that the Receiver has actually measured the states. This idea has been carried forward, and is the key to the proof of security that was developed 20 years later. In Protocol 5.1 we present a modern form of the protocol in the $\mathcal{F}_{\mathsf{so-com}}$-hybrid model, inspired by the version considered by Grilo et al. [Gri+20]. This is essentially the same as the one presented by Bennett et al., with the modifications they suggest.

Protocol 5.1 *BBCS92 in the $\mathcal{F}_{\mathsf{so-com}}$-hybrid model.*
The Sender S has strings $s_0, s_1 \in \{0,1\}^\ell$, the Receiver R has a choice bit $c \in \{0,1\}$.

1. **Preamble:** S sends $n > \ell$ BB84 states $|x^S\rangle_{\theta^S}$ prepared using random bits $x^S \leftarrow_R \{0,1\}^n$ and random choice of bases $\theta^S \leftarrow_R \{+,\times\}^n$.

   R measures the BB84 states in randomly chosen bases $\theta^R \leftarrow_R \{+,\times\}^n$, obtaining $x^R \in \{0,1\}^n$.

2. **Cut and choose:** R commits to the measurement outcomes and choice of bases, that is $\{\theta_i^R, x_i^R\}_{i\in[n]}$, using $\mathcal{F}_{\mathsf{so-com}}$.

S request to open a random subset $t \subseteq [n]$ of size $n/2$, and receives $\{\theta_i^R, x_i^R\}_{i \in T}$ from $\mathcal{F}_{\text{so-com}}$.

S aborts if for any $i \in t$ it has $\theta_i^S = \theta_i^R$ but $x_i^S \neq x_i^R$.

3. **Partition of index set:** S reveals $\theta_{\bar{t}}^S$, i.e. its bases for the unchecked locations. R partitions $\bar{t}$ into a subset of locations where it measured in the same bases as S, $I_c := \{i \in \bar{t} \colon \theta_i^S = \theta_i^R\}$, and the rest, $I_{1-c} := \bar{t} \setminus I_c$, and sends (I_0, I_1) to S.

4. **Secret transferring:** For $i = 0, 1$, S hides the string s_i using randomness extracted from $x_{I_i}^S$, via a two-universal hash function F, and sends $m_i := s_i \oplus f(x_{I_i}^S)$ for $i \in \{0,1\}$ and $f \leftarrow_R F$. R recovers $s_c = m_c \oplus f(x_{I_c}^R)$.

For completeness, we also include a diagram of the BBCS92 protocol in the same style as the diagrams for QKD from Section 4.3, see Figure 9.

Sender S $\leftarrow (s_0, s_1)$				**Receiver R** $\leftarrow c$
$x^S \leftarrow_R \{0,1\}^n$				
$\theta^S \leftarrow_R \{+,\times\}^n$	$\xrightarrow{\quad\quad\quad\quad\quad\quad\quad\quad}$	$\|x^S\rangle_{\theta^S}$		$\theta^R \leftarrow_R \{+,\times\}^n$
				Measure $\|x^S\rangle_{\theta^S}$ wrt. θ^R,
	$\xleftarrow{\text{ack}}$	$\mathcal{F}_{\text{so-com}}$	$\xleftarrow{x^R, \theta^R}$	obtaining x^R
$t \leftarrow_R \{t \subseteq [n] \colon \lvert t\rvert = n/2\}$	$\xrightarrow{t}$	$\mathcal{F}_{\text{so-com}}$		
	$\xleftarrow{x_t{}^R, \theta_t^R}$	$\mathcal{F}_{\text{so-com}}$		
Verify that $\forall i \in t$:				
$(\theta_i^S \neq \theta_i^R) \bigvee (x_i^R = x_i^S)$	$\xrightarrow{\quad\quad\quad\quad\quad\quad\quad\quad}$	θ^S		$I_c \leftarrow \{i \in \bar{t} \colon \theta_i^S = \theta_i^R\}$
				$I_{1-c} \leftarrow \bar{t} \setminus I_c$
	$\xleftarrow{\quad\quad\quad\quad\quad\quad\quad\quad}$	I_0, I_1		
$f \leftarrow_R F$, for $i \in \{0,1\}$:				
$m_i := s_i \oplus f(x_{I_i}^S)$	$\xrightarrow{\quad\quad\quad\quad\quad\quad\quad\quad}$	m_0, m_1, f		$s_c := m_c \oplus f(x_{I_c}^R)$
Output: $\perp$				Output: s_c

Figure 9: Illustration of the BBCS92 protocol. Using ack to indicate that $\mathcal{F}_{\text{so-com}}$. sends the expected acknowledgment indicating that it has received input to commit to.

That the protocol enjoys correctness follows immediately, since if both par-

ties follows the protocol, then $x^S_{I_c} = x^R_{I_c}$, so $f(x^S_{I_c}) = f(x^R_{I_c})$, and hence

$$m_c \oplus f(x^R_{I_c}) = s_c \oplus f(x^S_{I_c}) \oplus f(x^R_{I_c}) = s_c.$$

In the protocol n should be sufficiently large that we can use half the bits in the cut-and-chose step, and still be (sufficiently) sure that we get large enough sets I_0 and I_1 that a dishonest R cannot reasonably find $f(x^S_{I_{1-c}})$. Speaking of f, this should be a two-universal hash function, as defined in Definition 3.30, with range $\{0,1\}^\ell$. This serves as our randomness extractor, guaranteeing that even if a corrupted R knows some partial information about, or the value of, some bits in $x^S_{I_{c-1}}$, the value of $f(x^S_{I_{c-1}})$ is still close to uniformly random.

While the intuition behind the security of the protocol is simple, proving it formally is a harder challenge, which we will tackle in Section 5.2. For now, we note that intuitively a corrupted Sender cannot learn anything about c, since it only learns the values of θ^R_i for $i \in t$, and the corresponding bits are never used again. Hence, I_0 and I_1 are essentially (from the Senders perspective) just a random partition of $\bar{t}$. On the other hand, the intuition for a corrupted Receiver is that it has to measure the BB84 states from Alice before committing or have only negligible chance of committing to "correct" values for all n BB84 states — that is, for each state either $i \notin t$, $\theta^R_i \neq \theta^S_i$, or $x^R_i = x^S_i$. Since this is before R receives θ^S, the odds of R having measured "many more" than half of the BB84 states in the same bases as S is also negligible. Since measuring a state prepared in the computational basis with respect to the Hadamard basis produces a uniformly random bit, and vice versa, it is not possible for R to construct both I_c and I_{1-c} in such a way that f will not extract enough randomness to hide at least one of the messages.

5.2 Security of the BBCS92 protocol

We prove the security of the protocol against a malicious Sender and a malicious Receiver separately.

Theorem 5.2 *Security of BBCS92.*
The BBCS92 protocol implements $\mathcal{F}_{\mathsf{ot}}$ with quantum-computational security in the $\mathcal{F}_{\mathsf{so-com}}$-hybrid model.

Proof. This will follow immediately from Theorem 5.3 and Theorem 5.4. □

5.2.1 Security against a malicious Sender

The proof of security against a malicious Sender is an adaption of the proof by Grilo et al. [Gri+20] to our notation. This proof is itself an adaption of a proof by Unruh [Unr10].

Theorem 5.3
Protocol 5.1 implements $\mathcal{F}_{\mathsf{ot}}$ with quantum-computational security in the

$\mathcal{F}_{\text{so-com}}$-hybrid model against a malicious Sender.

Proof. In order to show that the protocol is secure in the sense of Definition 4.3 we first construct a simulator Sim for an arbitrary (fixed) adversary $\mathcal{A}$ as follows.

1. **Preamble:** Sim interacts with $\mathcal{A}$ as an honest Receiver would, except that Sim does not measure the received qubits yet. Denote the simulators randomly chosen bases as $\hat{\theta}^{\text{Sim}}$.

2. **Cut and choose:** Sim acts as $\mathcal{F}_{\text{so-com}}$ (and commits to a dummy message). Upon receiving t from $\mathcal{A}$, Sim measures the corresponding qubits with respect to $\hat{\theta}^{\text{Sim}}$, obtaining x_t^{Sim}. It then gives $\mathcal{A}$ $\{\theta_i^{\text{Sim}}, x_i^{\text{Sim}}\}_{i\in t}$, and aborts if $\mathcal{A}$ aborts.

3. **Partition of index set:** Denote by $\hat{\theta}^S$ the basis Sim receives from $\mathcal{A}$. Sim then measures the remaining qubits with respect to $\hat{\theta}^S$, and obtains $\hat{x}_{\bar{t}}^{\text{Sim}}$. Sim partitions $\bar{t}$ into two sets, I_0 and I_1. Sim randomly choose which one of them is the indices where $\hat{\theta}^S$ and $\hat{\theta}^{\text{Sim}}$ agrees. Sim sends I_0, I_1 to $\mathcal{A}$.

4. **Secret transferring:** Upon receiving m_0, m_1 from $\mathcal{A}$, Sim computes $s_i' := m_i \oplus f(\hat{x}_{I_i}^{\text{Sim}})$ for $i \in \{0,1\}$. Finally, Sim inputs (s_0', s_1') into $\mathcal{F}_{\text{ot}}$, and forwards whatever quantum (or classical) state that $\mathcal{A}$ outputs, since $\mathcal{A}$ has arbitrary output, and if the simulator outputs $\perp$ (as the honest Sender would) it would be trivial to distinguish between the real- and ideal-world.

See Figure 10 for a diagram of the ideal-world interactions, with the simulator as described above.

We claim that for the environment, the execution of BBCS92 with $\mathcal{A}$ is computationally indistinguishable from the execution of $\mathcal{F}_{\text{ot}}$ with Sim , i.e. that

$$\text{Out}_{\text{BBCS92},\mathcal{A},\text{R}} \overset{qc}{\approx} \text{Out}_{\mathcal{F}_{\text{ot}},\text{Sim},\hat{\text{R}}}. \tag{5.1}$$

To prove this we consider four hybrid arguments.

1. **Hybrid 1:** Sim_1 interacts with $\mathcal{A}$ in a protocol BBCS92_{H_1} that is exactly like BBCS92, but Sim_1 simulates honest R and $\mathcal{F}_{\text{so-com}}$, see Figure 11. It is therefore trivially the case that

$$\text{Out}_{\text{BBCS92},\mathcal{A},\text{R}} \overset{qc}{\approx} \text{Out}_{\text{BBCS92}_{H_1},\mathcal{A},\text{Sim}_1}. \tag{5.2}$$

2. **Hybrid 2:** See Figure 12. Sim_2 interacts with $\mathcal{A}$ in a protocol BBCS92_{H_2} that is like BBCS92_{H_1}, except that Sim_2 does not immediately measure the qubits, but instead measures at the following times and with respect to the following bases.

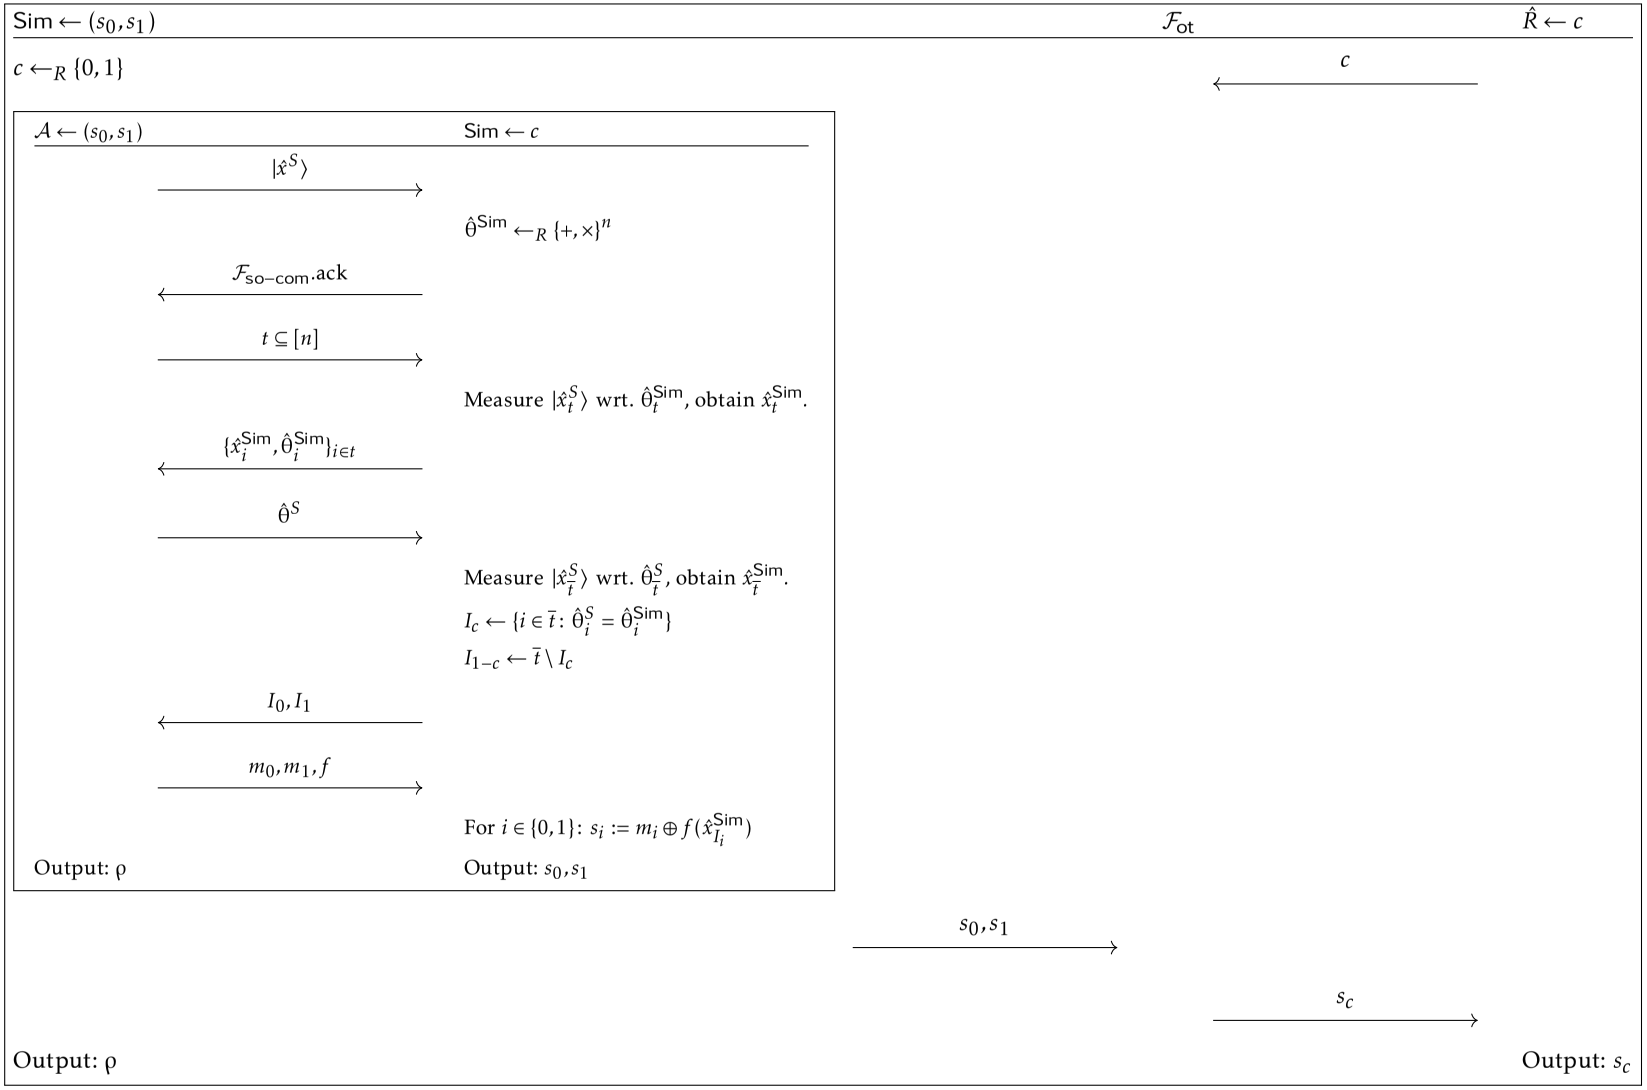

Figure 10: The ideal-world situation. We use $\mathcal{F}_{\text{so-com}}$.ack to indicate that Sim (acting in place of the ideal commitment functionality) sends the expected acknowledgment indicating that $\mathcal{F}_{\text{so-com}}$ has received input to commit to.

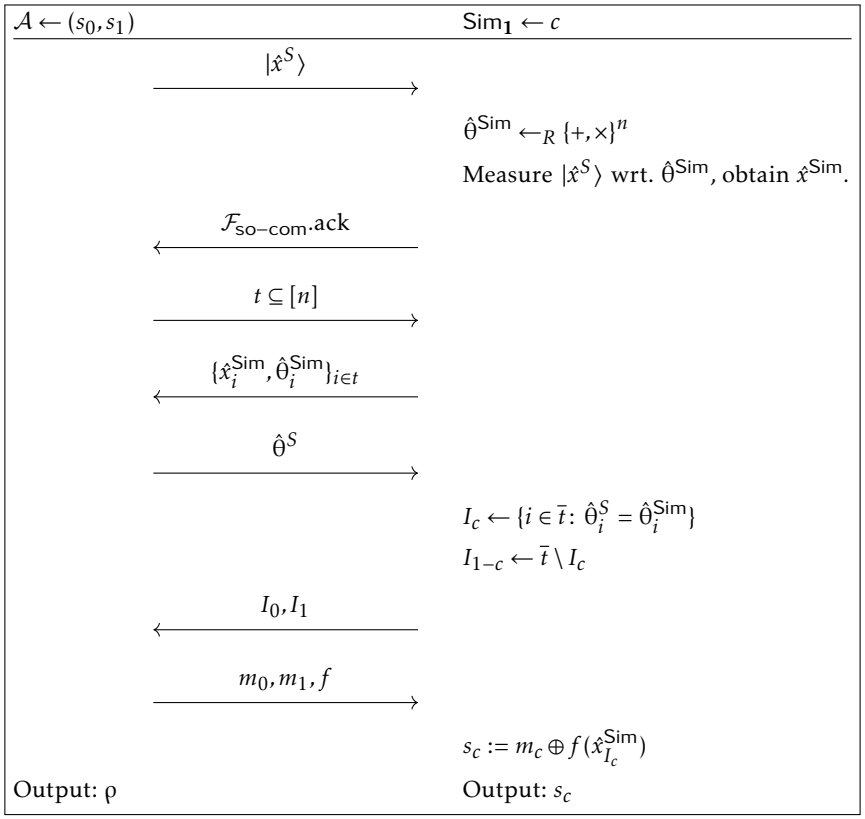

Figure 11: Hybrid argument H_1. Using $\mathcal{F}_{\mathsf{so-com}}$.ack to indicate that Sim_1 (acting in place of the ideal commitment functionality) sends the expected acknowledgment indicating that $\mathcal{F}_{\mathsf{so-com}}$ has received input to commit to.

(a) Sim_2 waits for the subset t before measuring the corresponding qubits with respect to θ^{Sim}, and then simulates opening $\mathcal{F}_{\mathsf{so-com}}$ with the used bases and measured outcomes.

(b) After receiving $\hat{\theta}^S$ from $\mathcal{A}$, Sim_2 measures the remaining qubits in θ^{Sim}, and then continues as in $\mathsf{BBCS92}_{H_1}$.

Since the only difference between the first hybrid and this one is that Sim_2 delays the measurements and operations on different subsystems commute, we have that

$$\mathrm{Out}_{\mathsf{BBCS92}_{H_1},\mathcal{A},\mathsf{Sim}_1} \overset{qc}{\approx} \mathrm{Out}_{\mathsf{BBCS92}_{H_2},\mathcal{A},\mathsf{Sim}_2}. \tag{5.3}$$

3. **Hybrid 3:** Sim_3 interacts with $\mathcal{A}$ in a protocol $\mathsf{BBCS92}_{H_3}$ that is like $\mathsf{BBCS92}_{H_2}$, except that Sim_3 measures the qubits in $|\hat{x}_{\bar{t}}^S\rangle$ with respect to $\hat{\theta}^S$ instead of θ^{Sim}. That is, as in Figure 12, but with $\hat{\theta}_{\bar{t}}^{\mathsf{Sim}}$ replaced with $\hat{\theta}_{\bar{t}}^S$. It follows that the only difference between hybrid 2 and 3 is the basis

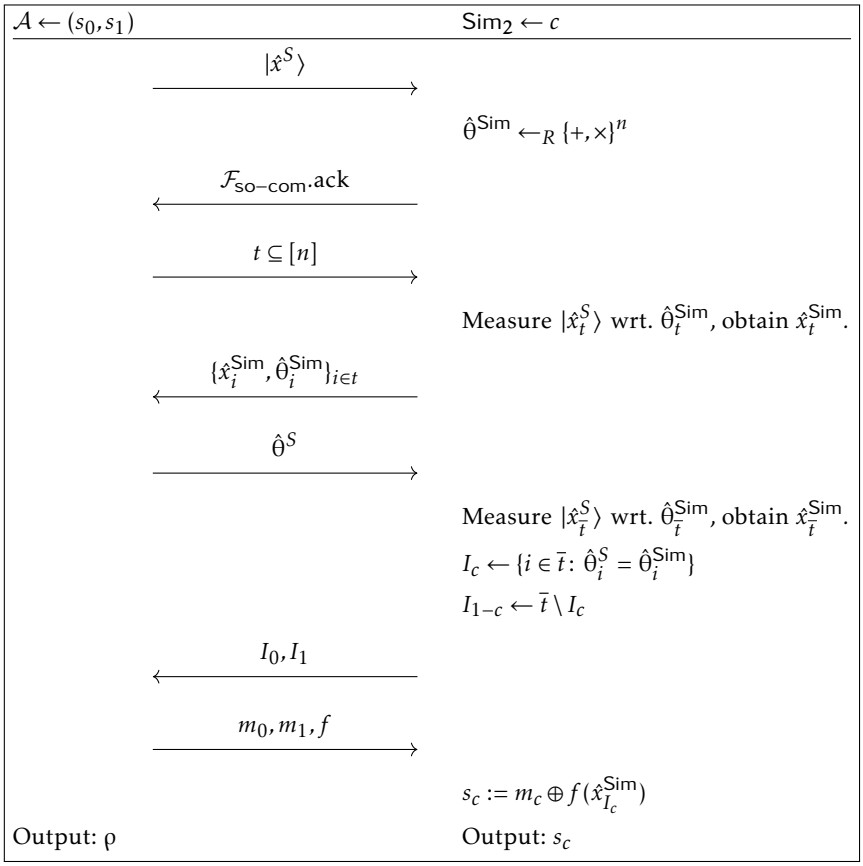

Figure 12: Hybrid argument H_2. Using $\mathcal{F}_{\mathsf{so-com}}.\mathsf{ack}$ to indicate that Sim_2 (acting in place of the ideal commitment functionality) sends the expected acknowledgment indicating that $\mathcal{F}_{\mathsf{so-com}}$ has received input to commit to.

used for measuring the qubits not in t, and since these are never revealed to $\mathcal{A}$, we see that

$$\mathrm{Out}_{\mathsf{BBCS92}_{H_2},\mathcal{A},\mathsf{Sim}_2} \overset{qc}{\approx} \mathrm{Out}_{\mathsf{BBCS92}_{H_3},\mathcal{A},\mathsf{Sim}_3}. \tag{5.4}$$

4. **Hybrid 4:** The simulator Sim interacts with $\mathcal{A}$, as described in the start of the proof. Sim performs the same operations as Sim_3 in $\mathsf{BBCS92}_{H_3}$ with respect to $\mathcal{A}$, and only uses information Sim_3 already had to extract both messages, that it then inputs into $\mathcal{F}_{\mathsf{ot}}$. It follows that

$$\mathrm{Out}_{\mathsf{BBCS92}_{H_3},\mathcal{A},\mathsf{Sim}_3} \overset{qc}{\approx} \mathrm{Out}_{\mathcal{F}_{\mathsf{ot}},\mathsf{Sim},\hat{\mathsf{R}}}. \tag{5.5}$$

At this point (5.1) follows immediately from transitivity of advantages being negligible, allowing us to combine equations (5.2), (5.3), (5.4), and (5.5). □

Note that in of Equation (5.2), (5.3), (5.4), and (5.5), it is the case that the joint outputs are statistically indistinguishable, so as a corollary to the proof

we note that the BBCS92 protocol implements $\mathcal{F}_{\mathsf{OT}}$ in the $\mathcal{F}_{\mathsf{so-com}}$ hybrid model with unconditional security against a malicious Sender.

5.2.2 Security against a malicious Receiver

Theorem 5.4
Protocol 5.1 implements $\mathcal{F}_{\mathsf{ot}}$ with quantum-computational security in the $\mathcal{F}_{\mathsf{so-com}}$-hybrid model against a malicious Receiver.

The idea behind the proof of security against a malicious Receiver is to use a standard proof technique, where we consider a variant of the protocol, where instead of generating quantum states based on random strings, the Sender generates EPR-pairs, sends one of the entangled qubits to the Receiver and then measures its own qubit at a later time. This allows us to consider the Senders actions in Step 2 as a sampling strategy and hence the results from Section 3.4 imply that with high probability the joint state is in a superposition of states with low Hamming weight within her subsystem. Further, this must therefore also be the case within $A_{I_{1-c}}$, for the $c \in \{0,1\}$ such that her choice of basis for I_{1-c} differs the most from the one the Receiver committed to. Then Corollary 3.42 implies that the string obtained from measuring this part of her subsystem in her basis has high min-entropy, and the privacy-amplification from Theorem 3.31 finishes the proof.

Our proof for security against a malicious Receiver is inspired by the proof by Grilo et al. [Gri+20], but modified to more closely fit the notion of security that is used both here and in their paper. The interpretation of the protocol as a sampling game is based on the work of Bouman and Fehr [BF10], but adapted from random-OT to $\binom{2}{1} - \mathsf{OT}$.

Proof of Theorem 5.4. We first state the EPR version of the BCCS92 protocol, which we will denote by $\mathsf{BBCS92}_{\mathsf{EPR}}$.

Protocol 5.5 $\mathsf{BBCS92}_{\mathsf{EPR}}$ *in the* $\mathcal{F}_{\mathsf{so-com}}$*-hybrid model.*
The Sender S has strings $s_0, s_1 \in \{0,1\}^\ell$, the Receiver R has a choice bit $c \in \{0,1\}$.

1. **Preamble:** S generates n EPR-pairs on the form $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, and sends R the second half of each of these EPR-pairs. S chooses random $\tilde{\theta}^S \leftarrow_R \{+,\times\}^n$, but does not measure her halves yet.

   R measures the received states in randomly chosen bases $\tilde{\theta}^R \leftarrow_R \{+,\times\}^n$, obtaining $\tilde{x}^R \in \{0,1\}^n$.

2. **Cut and choose:** As in BBCS92, R commits to the measurement outcomes and choice of bases, that is $\{\tilde{\theta}_i^R, \tilde{x}_i^R\}_{i\in[n]}$, using $\mathcal{F}_{\mathsf{so-com}}$.

   S request to open a random subset $t \subseteq [n]$ of size $n/2$, and receives $\{\tilde{\theta}_i^R, \tilde{x}_i^R\}_{i\in t}$ from $\mathcal{F}_{\mathsf{so-com}}$.

   Then, for each $i \in t$, S measures her i'th system with respect to $\tilde{\theta}_i^R$,

obtaining $\tilde{x}_i^S$. S aborts if for any $i \in t$, she has $\theta_i^S = \theta_i^R$ but $x_i^S \neq x_i^R$. After the checking is done S measures her remaining qubits in $\tilde{\theta}_{\bar{t}}^S$, obtaining $\tilde{x}_{\bar{t}}^S$.

3. **Partition of index set:** As in BBCS92.

4. **Secret transferring:** As in BBCS92.

As noted before the proof, the Receiver's view is identical in BBCS92 and $\mathsf{BBCS92}_{\mathsf{EPR}}$. Further, from the properties of EPR-states, it is immediate that an in-between version of this protocol where the Sender measures her states in Step 1, before sending anything to the Receiver, will give the Sender exactly the same output as in the BBCS92 protocol. Since the only difference between the in-between protocol and the $\mathsf{BBCS92}_{\mathsf{EPR}}$ protocol is that the Sender obtains her information at a later stage,[8] and since actions on different subsystems commute, it is immediate that for any adversary $\mathcal{A}$

$$\mathrm{Out}_{\mathsf{BBCS92},\mathsf{S},\mathcal{A}} = \mathrm{Out}_{\mathsf{BBCS92}_{\mathsf{EPR}},\mathsf{S},\mathcal{A}}.$$

The theorem now follows from showing that there is a simulator Sim such that

$$\mathrm{Out}_{\mathsf{BBCS92}_{\mathsf{EPR}},\mathsf{S},\mathcal{A}} \overset{qc}{\approx} \mathrm{Out}_{\mathcal{F}_{\mathsf{ot}},\hat{\mathsf{S}},\mathsf{Sim}}. \tag{5.6}$$

We construct the simulator Sim that takes the role of the Receiver in the ideal-world and uses the adversary $\mathcal{A}$ as a subroutine as follows.

1. (Preamble) Sim interacts with $\mathcal{A}$ acting as an honest Sender would in $\mathsf{BBCS92}_{\mathsf{EPR}}$.

2. (Cut and choose) Sim does the checking as the Sender would in $\mathsf{BBCS92}_{\mathsf{EPR}}$. Note that the commitments are simulated, so Sim sees all $\tilde{\theta}_i^R, \tilde{x}_i^R$, but only considers those that the Sender in the protocol would see. After checking, Sim measures her remaining subsystems, as the Sender would.

3. (Partition of index set) Sim receives (I_0, I_1) from $\mathcal{A}$.

4. (Secret transferring) The dummy Sender $\hat{S}$ inputs (s_0, s_1) to the ideal functionality $\mathcal{F}_{\mathsf{ot}}$. Sim inputs $c \in \{0,1\}$ to the ideal functionality, where c is chosen such that the bases that $\mathcal{A}$ has committed to and the bases that Sim has measured in, differs the most in the entries corresponding to I_{1-c}.

   Sim receives s_c from $\mathcal{F}_{\mathsf{ot}}$, and computes $m_c := s_c \oplus f(\tilde{x}_{I_c}^S)$ and randomly chooses $m_{1-c} \leftarrow_R \{0,1\}^\ell$. Send (m_0, m_1) to $\mathcal{A}$, and forwards whatever state $\mathcal{A}$ outputs.

Note that until the last step the simulation does not differ from the real-world interaction between the Sender and Receiver in $\mathsf{BBCS92}_{\mathsf{EPR}}$. In both executions

[8] And that the systems indexed by t are measured in $\tilde{\theta}^R$, but since they are only considered if $\tilde{\theta}_i^R = \tilde{\theta}_i^S$ anyway, this makes no difference.

they agree on their construction of m_c, but where $m_{1-c} = s_{1-c} \oplus f(\tilde{x}^S_{I_{1-c}})$ in the execution of $\mathsf{BBCS92}_{\mathsf{EPR}}$, it is chosen uniformly at random in the simulation. For the output to be the same, we must therefore show that after checking the min-entropy in the Senders halves of the EPR-pairs is high, even conditioned on the adversary's system. This will imply that f extracts ℓ uniformly random bits, so that the m_{1-c} is uniformly random in both executions, which will imply that (5.6) holds.

Towards this end we consider the state

$$|\phi_{AE}\rangle \in \mathcal{H}_{A_1} \otimes \cdots \otimes \mathcal{H}_{A_n} \otimes \mathcal{H}_E,$$

that is shared between the Sender and the adversary, after the adversary has committed to $\tilde{\theta}^R$ and $\tilde{x}^R$, but before the Sender chooses the test set t. Here $\mathcal{H}_{A_1} \otimes \cdots \otimes \mathcal{H}_{A_n}$ are the EPR-halves the Sender kept, and $\mathcal{H}_E$ is everything the adversary has, and possibly some additional states, making $|\phi_{AE}\rangle$ pure. Since $\mathcal{F}_{\mathsf{so-com}}$ is an ideal functionality, $\tilde{\theta}^R$ and $\tilde{x}^R$ are uniquely determined. To have our notation fit with Section 3.4, we assume that $\tilde{x}^R$ is the zero string, and that $\tilde{\theta}^R = (+,+,\ldots,+) \in \{+,\times\}^n$, i.e., the adversary has committed to zeroes, measured in the computational basis. However, as noted in Remark 3.38, this could be done for any string and for any choice of bases.

The essential observation is that the checking the Sender performs in Step 2 can be considered as a sampling strategy, applied to $|\phi_{AE}\rangle$, in order to check closeness of A to the all-zero state. Indeed, we see that this is exactly the sampling strategy considered in Example 2.21, adapted to test closeness to the all-zero state, by rejecting if $f(t, q_t, s) \neq 0$. To see this, note that the Sender choose a random subset $t \subseteq [n]$ of size $k = n/2$, measure A_t in the computational basis to obtain $q_t = \tilde{x}^S_t$ and either accept or reject based on testing a subset $s \subseteq t$, given by $s = \{i \in t \colon \tilde{\theta}^R_i = \tilde{\theta}^S_i\}$, which we observe is indeed uniformly random, since we chose $\tilde{\theta}^S$ uniformly at random.

If the test fails, i.e. there is some $i \in t$ such that $\tilde{\theta}^R_i = \tilde{\theta}^S_i$, but $\tilde{x}^R_i \neq \tilde{x}^S_i$, the Sender (respectively simulator) aborts, and Equations (5.6) trivially holds, so we assume the test does not fail. Interpreting this as a sampling game then allows us to conclude that at the end of the cut and choose phase, and for any $\delta > 0$, the joint state of $A_{\bar{t}}E$ is in a state $|\psi_{A_{\bar{t}}E}\rangle$ that is at most $\varepsilon^\delta_{\text{quant}}$ away from a superposition of states with relative Hamming weight at most δ within $A_{\bar{t}}$. We now assume that $|\psi_{A_{\bar{t}}E}\rangle$ acts exactly like this, and have here what we will refer to as our first error term, which by Theorem 3.39 and Example 2.21 is

$$\varepsilon^\delta_{\text{quant}} \leq \sqrt{\varepsilon^\delta_{\text{class}}} \leq \sqrt{6}\exp\left(-\frac{1}{100}k\delta^2\right).$$

At this point, Corollary 3.42 implies that if for $i = 0, 1$ we let X_i be the random variable for $\tilde{x}^S_{I_i}$, the string obtained by measuring the EPR halves in A_{I_i} in the Senders bases, then $\mathrm{H}_{\min}(\rho_{X_0X_1E}|E)$ is "large", however, we need that $\mathrm{H}_{\min}(\rho_{X_0X_1E}|X_cE)$ is "large" for some $c \in \{0,1\}$.

Towards this end, we note that interpreting $\tilde{\theta}^S$ as a bit string, where a 1 implies measuring in the Hadamard basis, we can consider $\omega(\tilde{\theta}^S_{\bar{t}})$ as the sample mean, and then apply Hoeffding's inequality to see that

$$\Pr\left[\left|\omega(\theta^S_{\bar{t}}) - \frac{1}{2}\right| \geq \varepsilon\right] \leq 2\exp\left(-2\varepsilon^2|t|\right) = 2\exp\left(-\varepsilon^2 n\right),$$

and note that this implies that also $\mathrm{wt}(\theta^S_{\bar{t}}) \geq (1/2-\varepsilon)n/2$, with at most probability $2\exp(-\varepsilon^2 n)$. We assume this bound holds, and have what we will refer to as our second error term,

$$\varepsilon_{\text{bases}} := 2\exp\left(-\varepsilon^2 n\right).$$

It follows that regardless of how the Receiver divides $\bar{t}$, there must be a $c \in \{0,1\}$, such that

$$\mathrm{wt}(\theta^S_{I_{1-c}}) \geq \frac{1}{2}\left(\frac{1}{2} - \varepsilon\right)\frac{n}{2},$$

since $\mathrm{wt}(\theta^S_{\bar{t}}) = \mathrm{wt}(\theta^S_{I_c}) + \mathrm{wt}(\theta^S_{I_{1-c}})$. For the rest of this proof fix this c. Note that if the Receiver acts honestly this c agrees with his input bit, and $\mathrm{wt}(\theta^S_{\bar{t}}) = \mathrm{wt}(\theta^S_{I_{1-c}})$, i.e. all the states in $\bar{t}$ they have measured in different bases are indexed by I_{1-c}.

Re-arranging the Sender's qubits, we can write $|\psi_{A_{\bar{t}}E}\rangle$ as $|\psi_{A_{I_0}A_{I_1}E}\rangle$ to ease notation. Since we assumed that we were not in any of the "error cases", $|\psi_{A_{\bar{t}}E}\rangle$ has relative Hamming weight at most δ within $A_{\bar{t}}$, or equivalently Hamming weight at most $n\delta/2$. Hence, it is also clear that $|\psi_{A_{I_0}A_{I_1}E}\rangle$ has Hamming weight at most $n\delta/2$ within $A_{I_{1-c}}$. Let X_0 and X_1 be the random variables for the measurement outcomes (as above), and let $\rho_{X_{1-c}X_cE}$ be the resulting classical-quantum state. We can then think of $\rho_{X_{1-c}X_cE}$ as being obtained by first measuring the subsystem $A_{I_{1-c}}$, obtaining $\rho_{X_{1-c}A_{I_c}E}$, and then measuring A_{I_c}. As earlier noted the order of measurements on different subsystems makes no difference.

If we now apply Corollary 3.42 to $\rho_{X_{1-c}A_{I_c}E}$ we obtain that

$$\begin{aligned}
\mathrm{H}_{\min}\left(\rho_{X_{1-c}A_{I_c}E}|A_{I_c}E\right) &\geq \mathrm{wt}(\theta^S_{I_{1-c}}) - h(\delta)|I_{1-c}| \\
&\geq \frac{1}{2}\left(\frac{1}{2} - \varepsilon\right)n - h(\delta)\frac{n}{2} \\
&= \frac{n}{2}\left(\frac{1}{4} - \frac{\varepsilon}{2} - h(\delta)\right).
\end{aligned}$$

As claimed in (3.5), measuring only destroys information, so this bound also holds for

$$\mathrm{H}_{\min}\left(\rho_{X_{1-c}X_cE}|X_cE\right).$$

Hence, the privacy amplification of Theorem 3.31 implies that the odds of telling apart $f(\tilde{x}^S_{I_{1-c}})$ from a uniformly randomly generated string is at most

$$\varepsilon_f := \frac{1}{2}\cdot 2^{-\frac{1}{2}\left(\mathrm{H}_{\min}\left(\rho_{X_{1-c}X_cE}|X_cE\right)-\ell\right)} \leq \frac{1}{2}\cdot 2^{-\frac{1}{2}\left(\frac{n}{2}\left(\frac{1}{4}-\frac{\varepsilon}{2}-h(\delta)\right)-\ell\right)}.$$

We can therefore conclude that, except with the following negligible probability, no distinguisher can tell the m_{1-c} in $\mathsf{BBCS92}_{\mathsf{EPR}}$ and the m_{1-c} from the simulation apart, and hence Equation (5.6) must hold; any non-negligible advantage there would directly translate to telling the m_{1-c}'s apart, contradicting this result.

Recalling that the trace distance indicates the greatest advantage one could have for distinguishing the states, it is clear that we can directly include our error terms in an upper bound for this, since doing so is equivalent to assuming we have probability 1 of telling the states apart if we end up in the "error case". Thus, we have the following upper bound on the error

$$\varepsilon^{\delta}_{\text{quant}} + \varepsilon_{\text{bases}} + \varepsilon_f \leq \sqrt{6}\exp\left(-\frac{1}{100}k\delta^2\right) + 2\exp\left(-\varepsilon^2 n\right) + \frac{1}{2}\cdot 2^{-\frac{1}{2}\left(\frac{n}{2}\left(\frac{1}{4}-\frac{\varepsilon}{2}-h(\delta)\right)-\ell\right)},$$

which is negligible, and hence finishes the proof. □

5.3 Commitments

The final piece that we need, to take the BBCS92 protocol from the $\mathcal{F}_{\mathsf{so-com}}$-hybrid model into MiniQCrypt, is a composable protocol for $\mathcal{F}_{\mathsf{so-com}}$ which uses only pq-OWFs. In this section we introduce and discuss some different candidates for such a protocol.

5.3.1 Commitments from post-quantum one-way functions

A natural first approach, is to try Naor's commitment scheme [Nao89], which is an elegant, conceptually simple, statistically binding, and computationally hiding commitment scheme that is contained in MiniQCrypt. Note that the scheme we use needs to be statistically binding for our sampling analysis to be applicable.

Example 5.6
Assume that $G\colon \{0,1\}^n \to \{0,1\}^{3n}$ is a cryptographically secure pseudorandom generator. A commitment to $b \in \{0,1\}$ is as follows. The Receiver sends a random string $y \leftarrow_R \{0,1\}^{3n}$ to the Committer, who picks a random seed $s \leftarrow_R \{0,1\}^n$, computes $c := G(s) \oplus b \cdot y$, and sends c to the Receiver. To open a commitment, the Committer sends b and s to the Receiver, who then verifies that $c = G(s) \oplus b \cdot y$.

The scheme is statistically binding, since for the Committer to cheat it must find a pair of s and s' such that $G(s') = G(s) \oplus y$. Since the range of G contains at most 2^n values, while y is chosen from 2^{3n} different values, there is only a $2^{2n}/2^{3n} = 2^{-n}$ chance that such s, s' even exists. The scheme is computationally hiding, since if the Receiver can distinguish a commitment to 0 from a commitment to 1, it can tell the output of G apart from the truly random $G(s) \oplus y$, breaking the security of G.

Since pseudorandom generators can be implemented from one-way functions, this commitment scheme is in MiniQCrypt. The problem with using this

commitment scheme in the BBCS92 protocol is that it is not at all clear how to argue that it is secure with respect to the notion of security from Definition 4.3, and hence the composition theorem is not applicable. We outline the problems, and how they would affect the BBCS92 protocol.

- First, it is not clear how to create an *efficient* simulator against a corrupted Committer (the Receiver in the BBCS92 protocol).
- Secondly — and more importantly — it is not at all clear how one would even go about creating a simulator (even ignoring efficiency) against a corrupted Receiver (Sender in the BBCS92 protocol), since we cannot rely on rewinding.[9]

Clearly a different approach is needed.

5.3.2 Commitments from LWE

In the first full proof of the security of the BBCS92 protocol, they investigated the properties that a commitment scheme should have for it to be composable and noted that such a scheme exists under the LWE-assumption [Dam+09]. In this section we take a closer look at their results. Before we start, we note that their proof is different from ours, and hence they do not need their commitments to be statistically binding.

In their proof, they have unconditional security against a malicious Sender (Receiver in the commitment scheme), which they want to preserve, and hence they need to use a statistical hiding and computational binding commitment scheme. On the other hand they also allow unbounded simulation against a malicious Sender. However, as we noted when discussing Naor's commitment scheme, the classical way of showing security against a malicious Committer is using rewinding, a technique that is not directly applicable in the quantum setting, so a different approach is needed. Towards this end, the authors suggest using a *dual-mode keyed commitment scheme*.

A dual-mode keyed commitment scheme is a commitment scheme, that has two different key-generation algorithms, $\mathcal{G}_H$ and $\mathcal{G}_B$. $\mathcal{G}_H$ (the *honest key-generation*) generates a public-key pk_H such that when using pk_H, the commitment scheme is statistically hiding and computationally binding. On the other hand, $\mathcal{G}_B$ (the *binding key-generation*) generates a public/private-key pair $(\mathsf{pk}_B, \mathsf{sk})$, such that when using pk_B, the commitment scheme is statistically binding, but only computationally hiding, and further the secret-key sk allows one to efficiently extract the values that are committed to. Finally, we require

[9] If the adversary has some internal quantum state that it measures at some point it may no longer be possible to rewind the adversary to the state it was in at a time before it measured. Eg., if the adversary measures a $|0\rangle$ qubit in the Hadamard basis to decide which of two actions it should take, the post-measurement state is either $|+\rangle$ or $|-\rangle$, and hence further measurements (i.e. after rewinding) of this qubit will always lead to it taking the same action.

pk_B and pk_H to be computationally indistinguishable, even against quantum adversaries.

We note that in the common reference string model (CRS-model), one assumes that the key is contained in the CRS, allowing a simulator to easily use pk_B instead of pk_H, when simulating against a malicious Committer. For simplicity, we work in the CRS-model, but one can take the scheme out of the CRS-model by using a simulatable coin-toss protocol to generate the CRS [DL09].

It is relatively straight forward to show that a dual-mode keyed commitment scheme that satisfies these assumptions, can be used for a secure implementation of $\mathcal{F}_{\mathsf{so-com}}$. We outline this below, but for details we refer to [Dam+09].

- For a corrupted Receiver, one simply uses unbound simulation and the fact that using pk_H the commitment scheme is only computationally binding, to open the commitments to the desired values. Looking back at the proof of security of the BBCS92 protocol, this allows the simulator to only start measuring qubits after receiving t from the malicious Sender.

- Against a corrupted Committer, the simulator generates $(\mathsf{pk}_B, \mathsf{sk})$, puts pk_B on the CRS, and extracts the values that the adversary commits to, using sk. As above, this clearly gives the simulator the power it needs in the proof of security for the BBCS92 protocol.

The problem with using this commitment scheme in the BBCS92 protocol, is that it assumes the existence of a public-key cryptosystem and is therefore not contained in MiniQCrypt.

5.3.3 Equivocable and extractable commitments

Ideally, we would like to somehow combine the commitment schemes from the previous two sections, into a scheme that has the properties that we need and is also in MiniQCrypt. Luckily for us, this was shown to be obtainable by Grilo et al. [Gri+20], and here we present the key points of their results.

For our commitment scheme to be composable with the BBCS92 protocol, we need it to have the following properties.

- **Equivocability:** When proving security against a corrupt Sender (Receiver in the commitment scheme), we need to be able to extract both the Senders messages, which we can then forward to the ideal functionality. To do this, we want the scheme to be such that the simulator can *equivocate* its commitments after receiving the challenge t from the adversary. If the simulator can do this, it can obtain the bases the adversary (claims to) have prepared the BB84 states in $A_{\bar{t}}$ with respect to before measuring these systems, and hence extract both strings.

- **Extractability:** When proving security against a corrupted Receiver (Committer in the commitment scheme), we need to be able to efficiently figure

out which of I_0 and I_1 contains the most indices where the bases of the adversary and the simulator match. To do this, the simulator needs to be able to efficiently *extract* $\hat{\theta}^R$. This task is non-trivial in the quantum setting, since we cannot use rewinding.

We split the path to archiving this into four steps.

Step 1 The first trick is to use a zero-knowledge proof[10] to achieve equivocable commitments, with inefficient simulation against corrupted Committers. To do this, we use Naor's commitment scheme, as described in Example 5.6, but instead of opening commitments as in Example 5.6, we now open a commitment by giving a zero-knowledge proof that the commitment is indeed a commitment to the claimed value. See Protocol 5.7 for details.

It follows immediately that a commitment can be equivocated to any value, simply by simulating the zero-knowledge protocol, a task that it is known that there are secure and efficiently simulatable protocols for in MiniQCrypt [Wat09].

Protocol 5.7 *Protocol 1 for* $\mathcal{F}_{\mathsf{so-com}}$.
The Committer has a string $m \in \{0,1\}^n$, the Receiver a set $t \subseteq [n]$.

- **Commitment phase:** The Receiver sends randomness ρ to the Committer. The Committer commits to m using Naor's commitment scheme and seeds s_1 to s_n. The Committer sends the produced commitments, $c_1, \ldots, c_n$, to the Receiver.
- **Decommitment phase:** When the Receiver requests to open t, the Committer sends m' and gives a zero-knowledge argument that $c_1, ..., c_n$ commits to m such that $m' = m_t$.

This protocol is trivially but inefficiently extractable, since Naor's commitment scheme is only computationally hiding. It is efficiently equivocable; simply commit to 0^n in the commitment phase, and in the decommitment phase send the m' one wishes to equivocate the message to and simulate the zero-knowledge argument.

Step 2 Observe that by plugging the above protocol (which securely implements $\mathcal{F}_{\mathsf{so-com}}$) into the BBCS92 protocol, we obtain a protocol that securely implements $\mathcal{F}_{\mathsf{OT}}$ in the plain model, with inefficient simulation against corrupted Receivers.

Additionally, it is straight forward to generalise Protocol 5.1 to a protocol for parallel OT. To perform k parallel transfers, we instead send $k \cdot n$ BB84 states, verifies them with a single instantiation of $\mathcal{F}_{\mathsf{so-com}}$, and then splits it so that

[10] We use "zero-knowledge proof" to be understood as a protocol where one party (the Prover) convinces some other party (the Verifier) that they know something, or that something was done according to some specification, without the Verifier learning anything (other than the fact that the Prover knows what it claims to know). For details, we refer to Chapter 21 of [Sma16].

the first n qubits are associated with the first OT, and so on. Plugging Protocol 5.7 into this modified version of the BBCS92 protocol, we obtain a protocol that securely implements $\mathcal{F}_{\text{p-OT}}$ in the plain model, with inefficient simulation against corrupted Receivers.

With a secure protocol for oblivious transfer, it follows that there are protocols for secure two-party computation of arbitrary classic functionalities [IPS08], again with inefficient simulation.

Step 3 Now, the key trick is to use this inefficiently simulatable protocol for $\mathcal{F}_{\text{p-OT}}$, together with zero-knowledge proofs, to construct a new protocol for $\mathcal{F}_{\text{so-com}}$, that *is* efficiently simulatable and have the desired properties. Towards this end, one defines the following functionality for *conditional disclosure of secrets*.

Definition 5.8 $\mathcal{F}_{\text{CDS}}$.
Let $\mathcal{L}$ be a NP language. $\mathcal{F}_{\text{CDS}}$ for the language $\mathcal{L}$ takes as input (x, m) from Alice and w from Bob. Alice receives no output, and Bob receives output (x, m') where m' is

$$m' := \begin{cases} m & \text{if } \mathcal{R}_{\mathcal{L}}(x, w) = 1, \\ \perp & \text{otherwise.} \end{cases}$$

I.e., this functionality is such that *if* Bob knows a witness for x, then Bob obtains m. It now follows from Step 2 that for any language there is a protocol that implements $\mathcal{F}_{\text{CDS}}$ with inefficient simulators.

We will be using the following language

$$\mathcal{L}_{\text{com}} := \{(c', b') \mid \exists r' \colon c' = \text{com}(b', r')\},$$

i.e. the language of pairs (c', b') such that c' is a valid commitment to b', using randomness r'. We will assume that the commitment scheme here is Naor's commitment scheme, instantiated with randomness y, which we write as com_y. For this language, a witness for (c', b') being in $\mathcal{L}_{\text{com}}$ is r', such that $c' = \text{com}_y(b', r') = G(r') \oplus b'y$. This implies that $\mathcal{L}_{\text{com}}$ is indeed a NP language.

We can now state our efficiently simulatable protocol for $\mathcal{F}_{\text{so-com}}$ in the $\mathcal{F}_{\text{CDS}}$-hybrid model.

Protocol 5.9 *Protocol 2 for* $\mathcal{F}_{\text{so-com}}$.
The Committer C has a string $m \in \{0,1\}^n$, the Receiver R a set $t \subseteq [n]$.

- **Commitment phase:**
  1. **Preamble:** C sends a random string ρ to R, and R sends a random string ρ^* to C.
  2. **Trapdoor setup:** R sends a Naor commitment $c = \text{com}_\rho(0, r)$ to C, and proves that c is a commitment to 0, i.e. that $((c, 0), r) \in \mathcal{R}_{\mathcal{L}_{\text{com}}}$,

using a zero-knowledge proof.

3. **CDS:** C acts as Alice in $\mathcal{F}_{CDS}$ with input $(x = (c, 1), m)$, and R as Bob with input $w = 0$.
4. **Commit:** C sends Naor commitments $c_i^* = \mathsf{com}_{\rho^*}(m, r^*)$ for $i \in [n]$ to R.

- **Decommitment phase:** R send t to C, and C sends m^* and proves via a zero-knowledge protocol that $c_1^*, \ldots, c_n^*$ is commitments to m such that $m_t = m^*$.

Note that if both parties follow the protocol this implements $\mathcal{F}_{\mathsf{so-com}}$ and, additionally, by the soundness of the zero-knowledge protocol, the Receiver cannot learn m, since c is a commitment to 0, and the Committer inputs $x = (c, 1)$. We observe that this protocol has the two properties that we need for the BBCS92 protocol.

- **Equivocability:** Just as in Protocol 5.7, the simulator can simulate the zero-knowledge proof in the decommitment phase.
- **Extractability:** By simulating the zero-knowledge proof in the trapdoor setup phase, the simulator can let c be a commitment to 1 instead of 0, and hence it can obtain m as output from $\mathcal{F}_{\mathsf{CDS}}$ by inputting the randomness r, that it used to generate c, instead of 0.

Clearly, both equivocation and extraction are efficient in the $\mathcal{F}_{\mathsf{CDS}}$-hybrid model, but the key observation is that all we need from the protocol that implements $\mathcal{F}_{\mathsf{CDS}}$ is input indistinguishability, and input indistinguishability is implied by inefficient simulation. To argue that input indistinguishability is sufficient, we note the following.

- A simulator against a corrupted Receiver proceeds as follows. In the first two steps it acts as an honest sender would, in the CDS it inputs the zero-vector into $\mathcal{F}_{\mathsf{CDS}}$, and similarly it also commits to the all-zero vector. In the decommitment phase, the simulator forwards t to $\mathcal{F}_{\mathsf{so-com}}$, receives m_t which it then sends to $\mathcal{A}$, and finally it equivocates the opening to m_t by simulates the zero-knowledge proof.

  That no distinguisher can tell apart this simulation and the real-world execution follows from the following. By the soundness of the zero-knowledge proof, c is a commitment to 0, and hence the output of $\mathcal{F}_{\mathsf{CDS}}$ is $(x, m' = \perp)$, and if the protocol implementing $\mathcal{F}_{\mathsf{CDS}}$ enjoys input indistinguishability, no adversary can tell this apart from the real execution. Indistinguishability for the last two points follows from Naor's scheme being secure and soundness of the last zero-knowledge proof.
- A simulator against a corrupted Committer uses a commitment to 1 instead of 0 in the trapdoor setup and simulates the zero-knowledge proof.

It can then receive m from $\mathcal{F}_{\mathsf{CDS}}$ by using the randomness used for the commitment as input. Finally, it inputs m into $\mathcal{F}_{\mathsf{so-com}}$, and runs the rest of the protocol as an honest Receiver would.

Indistinguishability once again follows from the soundness of the zero-knowledge protocol and input indistinguishability of the protocol that implements $\mathcal{F}_{\mathsf{CDS}}$.

Importantly, non of these arguments require us to simulate $\mathcal{F}_{\mathsf{CDS}}$, so the fact that the protocol that securely implements $\mathcal{F}_{\mathsf{CDS}}$ only has inefficient simulators is not a problem.

We note that there is still a problem with this protocol (we consider this in the next section), but this idea behind the protocol is sound and the problem can be fixed. This implies that if we take the protocol for $\mathcal{F}_{\mathsf{CDS}}$ with inefficient simulation, that was implied by Step 2 and plug it into Protocol 5.9, we obtain an efficiently simulatable protocol for $\mathcal{F}_{\mathsf{so-com}}$ in the plain model.

Step 4 The final step is to note that there are still a problem with protocol 5.9. In no way does it ensure that the m that the Committer inputs to $\mathcal{F}_{\mathsf{CDS}}$ is the same m that it commits to in the final step, and hence the protocol is not binding. Luckily, this can be fixed by using *verifiable conditional disclosure of secrets*, vCDS, which is a variation of CDS, where Alice outputs (x, m) and a proof π with the following two properties. First, the proof can be efficiently verified with respect to the transcript τ of the messages exchanged during the protocol, m, and x. Secondly, no malicious Alice can output (x, m, π) such that (1) (τ, x, m, π) can be verified, (2) Bob did not abort, and (3) Bob outputted m' that is inconsistent with the inputs (x, m) and w.

Using vCDS instead of CDS, we can add to the final step of the commitment phase that the Committer should prove that c^* is consistent with the input that the Committer used in the vCDS protocol, i.e. it should prove that (τ, x, m, π) can be verified, and that $c^* = \mathsf{com}_{\rho^*}(m, r^*)$. For details on creating a protocol for vCDS, and proving the security of this final version, we refer to [Gri+20].

Remark 5.10
As the closing remark of this section, I would like to give the opinion of the writer on this commitment scheme.

Treating the article, and the result that are obtained therein, as purely a proof of the existence of a protocol that securely implements $\mathcal{F}_{\mathsf{so-com}}$ in MiniQCrypt, I believe that the obtained result is correct and very interesting. However, I think that the way this is obtained is problematic. While it is common to use *ad hoc* additions to protocols to prevent cheating (eg. commitments in the BBCS92 protocol), the *ad hoc* use of conditional disclosure of secrets in the article is *purely* to allow a simulator to "cheat". This is particularly clear when one considered that the entirety of the idea behind

the initial commitment protocol (Protocol 5.7) is still present in the final protocol (Protocol 5.9), but now with the addition of what might appear as an unreasonable amount of work just to allow a simulator to efficiently extract what is committed to.

In essence, I find that this makes me question the practical value of the created protocol, in a deeper sense than just the massive overhead that it also suffers from does. Would any reasonable Committer want to participate in a protocol where such a large part of the protocol is not relevant for the task the protocol achieves, nor is it directly preventing anyone from cheating, but only making it so that a simulator can "cheat" against the Committer?

6 Concluding remarks

The main work of this thesis is a rigorous proof that the BBCS92 protocol implements oblivious transfer with security against quantum adversaries in the $\mathcal{F}_{\mathsf{so-com}}$-hybrid model. Towards proving this, we have introduced select topics from quantum mechanics, quantum information theory, quantum computation theory, and quantum cryptography . An essential tool for the proof of security is the notion of sampling games in both a classical and quantum setting, and a theorem relating the error of a given sampling strategy in the quantum setting to the error of the analogous game in the classical setting.

To take the protocol for oblivious transfer from the $\mathcal{F}_{\mathsf{so-com}}$-hybrid model and into the plain MiniQCrypt model, we presented the idea behind the construction of an equivocable and extractable commitment scheme, that relies only on post-quantum one-way functions. Finally, we argue what result that we believe one should conclude from the existence of this scheme, namely that it should be seen as a proof-of-existence, rather than a useful implementation of $\mathcal{F}_{\mathsf{so-com}}$.

Building on the results that has been shown in this thesis, one can show that there are secure protocols for the computation of any multi-party functionality in MiniQCrypt [IPS08; Unr10]. Additionally, with a protocol that securely realises any classical functionality, one can show that it is possible to securely evaluate any quantum circuit which takes as input a quantum state from each of the parties, and output a quantum state to each of them [DNS12]. This is the quantum equivalent of a non-interactive two-party functionality.

References

[BB84] Charles H. Bennett and Gilles Brassard. "Quantum cryptography: Public key distribution and coin tossing". In: *Proceedings of the International Conference on Computers, Systems & Signal Processing* (1984), pp. 175–179.

[Ben+91] Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Marie-Hélène Skubiszewska. "Practical Quantum Oblivious Transfer". In: *Advances in Cryptology - CRYPTO '91*. Ed. by Joan Feigenbaum. Vol. 576. Lecture Notes in Computer Science. Springer, 1991, pp. 351–366.

[BF10] Niek J. Bouman and Serge Fehr. "Sampling in a Quantum Population, and Applications". In: *Advances in Cryptology - CRYPTO 2010*. Ed. by Tal Rabin. Vol. 6223. Lecture Notes in Computer Science. The full version with appendicies is avalible on arXiv. Springer, 2010, pp. 724–741.

[Bru+14] Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. "Bell nonlocality". In: *Rev. Mod. Phys.* 86 (2014), pp. 419–478.

[CDN15] Ronald Cramer, Ivan Bjerre Damgård, and Jesper Buus Nielsen. *Secure Multiparty Computation and Secret Sharing*. Cambridge University Press, 2015.

[Cle86] Richard Cleve. "Limits on the Security of Coin Flips When Half the Processors Are Faulty". In: *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*. STOC '86. Association for Computing Machinery, 1986, 364–369.

[Cré87] Claude Crépeau. "Equivalence Between Two Flavours of Oblivious Transfers". In: *Advances in Cryptology - CRYPTO '87*. Ed. by Carl Pomerance. Vol. 293. Lecture Notes in Computer Science. Springer, 1987, pp. 350–354.

[Dam+09] Ivan Damgård, Serge Fehr, Carolin Lunemann, Louis Salvail, and Christian Schaffner. "Improving the Security of Quantum Protocols via Commit-and-Open". In: *Advances in Cryptology - CRYPTO 2009*. Ed. by Shai Halevi. Vol. 5677. Lecture Notes in Computer Science. Springer, 2009, pp. 408–427.

[Dia+16] Eleni Diamanti, Hoi-Kwong Lo, Bing Qi, and Zhiliang Yuan. "Practical challenges in quantum key distribution". In: *npj Quantum Information* 2 (2016).

[DL09] Ivan Damgård and Carolin Lunemann. "Quantum-Secure Coin-Flipping and Applications". In: *Advances in Cryptology - ASIACRYPT 2009*. Ed. by Mitsuru Matsui. Vol. 5912. Lecture Notes in Computer Science. Springer, 2009, pp. 52–69.

[DNS12] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. "Actively Secure Two-Party Evaluation of Any Quantum Operation". In: *Advances in Cryptology - CRYPTO 2012*. Ed. by Reihaneh Safavi-Naini and Ran Canetti. Vol. 7417. Lecture Notes in Computer Science. Springer, 2012, pp. 794–811.

[Gol04] Oded Goldreich. *Foundations of Cryptography: Volume II - Basic Applications*. Cambridge University Press, 2004.

[Gra21] Federico Grasselli. *Quantum Cryptography: From Key Distribution to Conference Key Agreement*. Springer Berlin Heidelberg, 2021.

[Gri+20] Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. *Oblivious Transfer is in MiniQCrypt*. Accepted for EUROCRYPT 2021. 2020. arXiv: 2011.14980.

[Hid19] Jack D. Hidary. *Quantum Computing: An Applied Approach*. Springer, 2019.

[Hoe63] Wassily Hoeffding. "Probability Inequalities for Sums of Bounded Random Variables". In: *Journal of the American Statistical Association* 58 (1963), pp. 13–30.

[HSS15] Sean Hallgren, Adam Smith, and Fang Song. "Classical cryptographic protocols in a quantum world". In: *International Journal of Quantum Information* 13 (2015).

[IK97] Yuval Ishai and Eyal Kushilevitz. "Private Simultaneous Messages Protocols with Applications". In: *Fifth Israel Symposium on Theory of Computing and Systems*. IEEE Computer Society, 1997, pp. 174–184.

[Imp95] Russell Impagliazzo. "A Personal View of Average-Case Complexity". In: *Proceedings of the Tenth Annual Structure in Complexity Theory Conference*. IEEE Computer Society, 1995, pp. 134–147.

[IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. "Founding Cryptography on Oblivious Transfer - Efficiently". In: *Advances in Cryptology - CRYPTO 2008*. Ed. by David A. Wagner. Vol. 5157. Lecture Notes in Computer Science. Springer, 2008, pp. 572–591.

[JL03] Richard Jozsa and Noah Linden. "On the role of entanglement in quantum-computational speed-up". In: *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 459 (2003), pp. 2011–2032.

[Lo97] Hoi-Kwong Lo. "Insecurity of quantum secure computations". In: *Phys. Rev. A* 56 (1997), pp. 1154–1162.

[Mac09] Barbara D. MacCluer. *Elementary Functional Analysis*. Springer, 2009.

[Nao89] Moni Naor. "Bit Commitment Using Pseudo-Randomness". In: *Advances in Cryptology - CRYPTO '89*. Ed. by Gilles Brassard. Vol. 435. Lecture Notes in Computer Science. Springer, 1989, pp. 128–136.

[NC02] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2002.

[Reg05] Oded Regev. "On Lattices, Learning with Errors, Random Linear Codes, and Cryptography". In: *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*. STOC '05. Association for Computing Machinery, 2005, pp. 84–93.

[Ren05] Renato Renner. "Security of Quantum Key Distribution". PhD thesis. Swiss Federal Institute of Technology, Zurich, 2005.

[RK05] Renato Renner and Robert König. "Universally Composable Privacy Amplification Against Quantum Adversaries". In: *Theory of Cryptography*. Ed. by Joe Kilian. Vol. 3378. Lecture Notes in Computer Science. Springer, 2005, pp. 407–425.

[Sch17] René L. Schilling. *Measures, Integrals and Martingales*. 2nd ed. Cambridge University Press, 2017.

[Sma16] Nigel P. Smart. *Cryptography Made Simple*. 1st ed. Springer, 2016.

[Unr10] Dominique Unruh. "Universally Composable Quantum Multi-party Computation". In: *Advances in Cryptology - EUROCRYPT 2010*. Ed. by Henri Gilbert. Vol. 6110. Lecture Notes in Computer Science. Springer, 2010, pp. 486–505.

[Wat09] John Watrous. "Zero-knowledge against quantum attacks". In: *SIAM Journal on Computing* 39.1 (2009), pp. 25–58.

[WW06] Stefan Wolf and Jürg Wullschleger. "Oblivious Transfer Is Symmetric". In: *Advances in Cryptology - EUROCRYPT 2006*. Ed. by Serge Vaudenay. Vol. 4004. Lecture Notes in Computer Science. Springer, 2006, pp. 222–232.