Quantum Algorithms

Andreas Klappenecker *Texas A&M University*

Contents

Preface

Quantum computing provides a fresh perspective on information processing. Some quantum algorithms have the promise to provide an exponential speed-up over classical deterministic and randomized algorithms. This explains the massive worldwide efforts to build a viable quantum computer. However, this is certainly not the only motivation to study the subject matter. Quantum computing has serious repercussions on classical computing. For instance, some efficient algorithms for hard problems have been obtained by "de-quantizing" quantum algorithms.

These lecture notes provide a rapid introduction to the main ideas behind quantum algorithms. The subject matter is not *difficult*, but dramatically *different* from its classical counterpart. We provide numerous simple exercises that are designed to ease the transition into the quantum realm. Solving the exercises will help the reader to gain an active working knowledge.

Our approach is largely based on the quantum circuit model, which is easy to understand. This model abstracts from the nature and the dynamics of the physical system realizing the quantum computer. The advantage of this approach is that within an extremely short period of time it will be possible to cover interesting algorithms.

The course requires some background in linear algebra. The books *Linear Algebra* by Serge Lang and *Linear Algebra Done Right* by Sheldon Axler are excellent sources to review such material.

Please note that this is a preliminary draft. The lecture notes are incomplete, and all parts are subject to change. The material should be read in conjuction with the book *Quantum Computation and Quantum Information* by Nielsen and Chuang. You can consult the books *Classical and Quantum Computation* by Kitaev, Shen, and Vyalyi, and *Quantum Computing* by Hirvensalo for further information. Both are accessible for readers with a background in Computer Science. The most important additional resource, however, is the quantum physics archive `www.arxiv.org`, where you can find recent preprints.

If you read these notes, then you accept the following contract: You agree to communicate all errors to me. If you do not want to burden yourself with this task, then do not read any further.

Andreas Klappenecker
College Station, Texas

Chapter 1

Prolegomena

The predominant computational models of the last fifty years are all based on the notion of a bit, a representation of two alternatives, 0 or 1. The technological convenience to store and transmit information in such a form is evident. From that point of view, it might be surprising that the very concept of a bit is challenged by a contemporary computational theory.

The basic unit of information in the quantum computation model is a **quantum bit**. We motivate and illustrate this concept by a simple example, before developing the theory in a more axiomatic way. Hopefully, it will become apparent that the notion of a quantum bit is as natural as the notion of a bit.

We begin our journey by considering a source of monochromatic light. It is possible to polarize light such that it has an electric field which oscillates horizontally

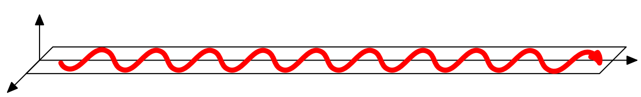

or vertically

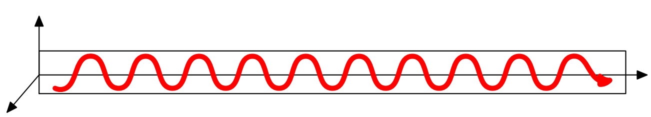

Suppose that we dim our monochromatic light source so that it emits a single photon at a time. The photon has the polarization property as well; hence, we can encode a classical bit by polarizing the photon either horizontally or vertically. We denote the two alternatives by $|\leftrightarrow\rangle$ and $|\updownarrow\rangle$ instead of 0 and 1, for reasons that will become apparent shortly.

We can distinguish the two different states by sending the photon through a calcite crystal. A horizontally polarized photon will pass straight through, whereas a vertically polarized photon will be deflected:

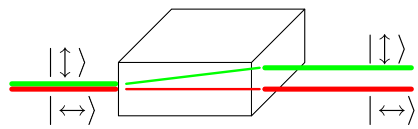

You can find such calcite crystrals for example in museum shops. The above effect can be verified by shining with a laser pointer through such crystals.

Rotations. There is no particular reason to single out horizontally and vertically polarized photons. For instance, we could have rotated the polarization planes by an angle θ. A photon can encode a bit just as well using the two rotated polarization directions.

An interesting aspect emerges by comparing the two representations. The rotated basis can be expressed in terms of the horizontal and vertical polarized states, $|\leftrightarrow\rangle$ and $|\updownarrow\rangle$, as follows:

$$\begin{aligned} |0_\theta\rangle &= \cos\theta|\leftrightarrow\rangle + \sin\theta|\updownarrow\rangle, \\ |1_\theta\rangle &= -\sin\theta|\leftrightarrow\rangle + \cos\theta|\updownarrow\rangle. \end{aligned}$$

The geometrical meaning of these formulas become immediately apparent if we associate the horizontally polarized state $|\leftrightarrow\rangle$ with the column vector $(1,0)^t$ and the vertically polarized state $|\updownarrow\rangle$ with $(0,1)^t$. We do not want to get into the interesting details of the underlying physics. For the moment, we will content ourselves by discussing a single consequence of the above formula that will lead us to an interesting application.

Suppose that we repeat the previous experiment with the calcite crystal, but this time we send a beam of photons that are all in the state $|0_\theta\rangle$. Assume that the calcite crystal is still aligned such that horizontally and vertically polarized photons can be perfectly distinguished. The first observation is that the photons emerging from the crystal all have either the polarization $|\leftrightarrow\rangle$ or $|\updownarrow\rangle$. Even more interesting is the fact that the emerging photon will be in the state $|\leftrightarrow\rangle$ with probability $\cos^2\theta$, and in the state $|\updownarrow\rangle$ with probability $\sin^2\theta$.

Exercise 1.1 *What angle θ do you have to choose such that the calcite crystal realizes a fair coin flip when presented with a photon in the state $|0_\theta\rangle$? Assume that the axis of the crystal is aligned such that it discriminates perfectly between $|\leftrightarrow\rangle$ and $|\updownarrow\rangle$.*

Polarization States. We can express more generally all other forms of polarization in terms of a linear combination $a|{\leftrightarrow}\rangle + b|{\updownarrow}\rangle$ of the horizontal and vertical polarization states, where a and b are complex numbers that satisfy $|a|^2 + |b|^2 = 1$.

There is nothing strange about the fact that we can have complex numbers as coefficients. For instance, a right-hand circularly polarized photon is a photon in the state

$$\frac{1}{\sqrt{2}}(|{\leftrightarrow}\rangle + i\,|{\updownarrow}\rangle).$$

This is the quantum analogue of right-hand circularly polarized light, which is the superposition of a horizontally oscillating electric field and a vertically oscillating electric field that are 90° out of phase. The coefficient i accounts for this difference in phase.

Exercise 1.2 *Explain why the term circularly polarized light is appropriate. Which polarization state of the photon would correspond to left-hand circular polarization?*

The polarization of a photon is an instance of a quantum bit, as you might have guessed already. A quantum bit has two clearly distinguishable states, in our case $|{\leftrightarrow}\rangle$ and $|{\updownarrow}\rangle$. The quantum bit can exist in superpositions of these two states, such as $a|{\leftrightarrow}\rangle + b|{\updownarrow}\rangle$. This is just a strange way to express a two-dimensional nonzero vector $(a, b)^t \in \mathbf{C}^2$. We can assure the reader that the notation will turn out to be immensely convenient.

It is not possible to extract the coefficient a and b of a superposition state $a|{\leftrightarrow}\rangle + b|{\updownarrow}\rangle$. If we want to learn something about the state, then we can send the photon, for instance, through the calcite crystal. If the crystal is aligned as before, then the outcome will be $|{\leftrightarrow}\rangle$ with probability $|a|^2$, and $|{\updownarrow}\rangle$ with probability $|b|^2$. The superposition collapses when we perform such a measurement. We can take advantage of this fact to realize a protocol for the secure distribution of keys.

Key Distribution. Establishing a common secret between two parties is an important cryptographical primitive. If Alice wants to send a confidential message to Bob over a public channel, then they can use encryption to prevent an eavesdropper from reading the message. If they use some standard block cipher such as AES or 3DES, then they need to have a common key. Public key cryptosystems, such as RSA, provide methods that can establish such a common secret.

There is a problem, however. An eavesdropper can silently copy all messages used to establish the key, and the encrypted message. The eavesdropper might not be able to take immediate advantage of the copied material. Nevertheless, she might be able to break the system later, and decipher the message.

In 1984, Bennett and Brassard introduced a protocol that allows to exchange a key securely. The protocol takes advantage of quantum mechanics to ensure that eavesdropping during the key exchange phase will not go unnoticed. This is an example of a property that cannot be guaranteed by any protocol that is based on classical physics.

Alice uses four different polarization states of photons in this protocol. The horizontally and vertically polarized states, $|\leftrightarrow\rangle$ and $|\updownarrow\rangle$, and a basis that is obtained by a 45° degree rotation,

$$|\swarrow\!\!\!\nearrow\rangle = \frac{1}{\sqrt{2}}|\leftrightarrow\rangle + \frac{1}{\sqrt{2}}|\updownarrow\rangle \qquad \text{and} \qquad |\nwarrow\!\!\!\searrow\rangle = \frac{1}{\sqrt{2}}|\leftrightarrow\rangle - \frac{1}{\sqrt{2}}|\updownarrow\rangle.$$

A classical bit can be encoded either by the alternatives $\boxplus = \{|\leftrightarrow\rangle, |\updownarrow\rangle\}$ or by $\boxtimes = \{|\swarrow\!\!\!\nearrow\rangle, |\nwarrow\!\!\!\searrow\rangle\}$. Alice and Bob agree on the following representation:

basis	encoding
$\boxplus$	$0 \cong |\leftrightarrow\rangle, 1 \cong |\updownarrow\rangle$
$\boxtimes$	$0 \cong |\swarrow\!\!\!\nearrow\rangle, 1 \cong |\nwarrow\!\!\!\searrow\rangle$

Bob can use a calcite crystal to measure a photon sent by Alice. He selects between two different alignments of the crystal. The alignment $\boxplus$ allows Bob to perfectly discriminate between $|\leftrightarrow\rangle$and $|\updownarrow\rangle$, and the alignment $\boxtimes$ to perfectly discriminate between $|\swarrow\!\!\!\nearrow\rangle$ and $|\nwarrow\!\!\!\searrow\rangle$. The second alignment is obtained from the first by rotating the calcite crystal by 45°. The following observation is crucial to the protocol:

Observation *If Alice sends a bit choosing one encoding, but Bob has aligned his crystal to measure the other, then he will decode* 0 *with probability* 1/2 *and* 1 *with probability* 1/2.

Protocol BB84. The goal of this protocol is to establish a common secret of n bits between Alice and Bob.

1) Alice chooses a data string s of $(4+\delta)n$ bits that are independently selected uniformly at random.

2) Alice chooses a string b of $(4+\delta)n$ symbols over the alphabet $\{\boxplus, \boxtimes\}$ that are independently selected uniformly at random.

3) For all $k \in \{1, \ldots, (4+\delta)n\}$, Alice sends the data bit s_k encoded in the basis b_k to Bob.

4) Bob selects for each incoming photon a basis from the set $\{\boxplus, \boxtimes\}$, independently and uniformly at random, and measures the photon in that basis. He records the basis that he has chosen and the measurement outcome.

5) Alice publicly announces the string b.

6) Alice and Bob discard all bits from s where Bob measured in the wrong basis. With high probability, there are at least $2n$ bits left. They repeat the protocol if that is not the case. They keep $2n$ bits.

7) Alice selects n bits from this string and announces the position and value of these bits. Bob compares the value of these n check bits with the values of the bits that he has measured. If more than an acceptable number disagree, then they abort the protocol.

8) Alice and Bob extract from the remaining n common bits a common key using information reconciliation and privacy amplification methods.

The purpose of the last step is to take into account that the state of some photons might have been disturbed by some imperfection of the communication channel. We will ignore the technical details of this last step for the time being. The following example illustrates the protocol:

s	0	1	1	0	0	0	1	1	0	1	0	0	1	1	1	0
b	$\boxtimes$	$\boxplus$	$\boxtimes$	$\boxtimes$	$\boxplus$	$\boxplus$	$\boxplus$	$\boxtimes$	$\boxtimes$	$\boxplus$	$\boxplus$	$\boxtimes$	$\boxplus$	$\boxtimes$	$\boxtimes$	$\boxplus$
polarization	$\lvert\nearrow\!\!\!\!\swarrow\rangle$	$\lvert\updownarrow\rangle$	$\lvert\nwarrow\!\!\!\!\searrow\rangle$	$\lvert\nearrow\!\!\!\!\swarrow\rangle$	$\lvert\leftrightarrow\rangle$	$\lvert\leftrightarrow\rangle$	$\lvert\updownarrow\rangle$	$\lvert\nwarrow\!\!\!\!\searrow\rangle$	$\lvert\nearrow\!\!\!\!\swarrow\rangle$	$\lvert\updownarrow\rangle$	$\lvert\leftrightarrow\rangle$	$\lvert\nearrow\!\!\!\!\swarrow\rangle$	$\lvert\updownarrow\rangle$	$\lvert\nwarrow\!\!\!\!\searrow\rangle$	$\lvert\nwarrow\!\!\!\!\searrow\rangle$	$\lvert\leftrightarrow\rangle$
Bob's basis	$\boxtimes$	$\boxplus$	$\boxplus$	$\boxplus$	$\boxtimes$	$\boxtimes$	$\boxplus$	$\boxtimes$	$\boxplus$	$\boxplus$	$\boxtimes$	$\boxtimes$	$\boxplus$	$\boxplus$	$\boxtimes$	$\boxtimes$
Detected bit	0	1	0	1	1	0	1	1	0	1	1	0	1	0	1	0
Correct basis?	✓	✓					✓	✓		✓		✓	✓		✓	
Check bits	0	1						1							1	
$\Rightarrow$ no eavesdropper																
Common secret							1			1		0	1			

What makes this protocol secure? All operations in quantum mechanics are linear. One consequence of this fact is that one cannot copy an unknown quantum state without disturbing the state. So an eavesdropper will not go unnoticed when she is trying to copy the bits. There is of course much more to say about this protocol, but we do not want to get into too many details in this introductory chapter. Devices realizing this protocol are already commercially available from a company in the USA and from a company in Europe.

Quo Vadis? The polarization state of a photon is just one potential way to store information. Other quantum mechanical systems can serve the same purpose. In fact, a staggering number of different quantum systems have been proposed for quantum information processing. Each system has some advantages and some disadvantages. We will not be concerned with the details of such proposals. We focus instead on properties that almost all of these proposals try to accomplish.

The key distribution protocol by Bennett and Brassard exemplified a few aspects of quantum information processing, but not all of them, not even close! The most interesting aspects emerge from the combination of several quantum systems. We will discuss operations that allow to manipulate such quantum memories. And we show how to harness various quantum features to obtain beautiful algorithms.

Chapter 2

Quantum Circuits

Quantum computing can be based on various different computational models. The most accessible one is the quantum circuit model, which specifies a sequence of operations that manipulate the state of the quantum computer at discrete time steps. The basic rules of this model are surprisingly simple. This chapter introduces the basic properties of quantum states, quantum gates, and measurements.

§1 Quantum States

A bit has two distinguishable states, denoted by 0 and 1. A classical computer manipulates a set of bits, which form the memory of the computer. The memory of a quantum computer is based in a similar way on the notion of a **quantum bit**, **qubit** for short. A qubit has two clearly distinguishable states, denoted by $|0\rangle$ and $|1\rangle$. The possible states of a qubit are not exhausted by these two possibilities. In general, the state of a qubit is of the form $a|0\rangle + b|1\rangle$, where a and b are complex numbers satisfying $|a|^2 + |b|^2 = 1$.

The states $|0\rangle$ and $|1\rangle$ should be understood as basis vectors of a complex two-dimensional vector space. We can associate with these states the basis vectors

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \tag{2.1}$$

The state $a|0\rangle + b|1\rangle$ is a linear combination of these two basis vectors, and is represented by the vector $(a, b)^t$. The operations of the quantum computer manipulate these vectors by linear transformations or by measurements.

The value of a quantum bit is always 0 or 1, never anything else. If a qubit is in the state $a|0\rangle + b|1\rangle$, then this means that the value 0 is observed with probability $|a|^2$, and the value 1 with probability $|b|^2$. A measurement in the computational basis returns the value 0 or 1 according to this rule, and sets the qubit to the state $|0\rangle$ or $|1\rangle$, respectively. A consequence is that if the measurement is repeated, then it will return the same value.

It is easy to construct a state that yields a fair coin-flip. Choose the state $\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$. Then 0 and 1 are both observed with probability $(1/\sqrt{2})^2 = 1/2$. The resulting state after the measurement is $|0\rangle$, if the measurement result was 0, and $|1\rangle$ otherwise.

Exercise 2.1 *Assume that a qubit is in the state* $\frac{1}{\sqrt{10}}|0\rangle + \frac{3}{\sqrt{10}}|1\rangle$. *What is the probability to observe* 0, *or* 1, *respectively?*

Exercise 2.2 *Assume that a qubit is in the state* $\frac{i}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle$. *What is the probability to observe* 0, *or* 1, *respectively?*

A memory consisting of n quantum bits has 2^n basis states, which are denoted by $|0\cdots 00\rangle, |0\cdots 01\rangle, |0\cdots 10\rangle, \ldots, |1\cdots 11\rangle$. The state of the memory is a linear combination of these basis states. Denote by $\mathbf{F}_2$ the finite field with two elements 0 and 1. An arbitrary state of the memory is of the form

$$\sum_{k\in\mathbf{F}_2^n} a_k|k\rangle, \quad \text{with} \quad \sum |a_k|^2 = 1.$$

If we read out the memory by a measurement in the computational basis, then we will observe the result k, a string of n bits, with probability $|a_k|^2$. The scalar coefficients a_k are called **probability amplitudes** or, simply, **amplitudes**.

Exercise 2.3 *What is the probability of observing* 11, *if the memory is in the state* $\frac{1}{2}|00\rangle - \frac{1}{2}|10\rangle + \frac{i}{\sqrt{2}}|11\rangle$? *In what state is the memory once we have observed* 11?

Exercise 2.4 *Describe all possible states of a system of two quantum bits such that a measurement in the computational basis yields* 00 *with probability* 1/2, *and the results* 01 *and* 11 *both with probability 1/4.*

Any quantum system with at least two different basis states can basically store a quantum bit, and finding appropriate storage media for a quantum computer is a very active area of current research. The linear combination of basis states reflects the superposition principle of quantum mechanics. It should be noted that only the measurement process introduces randomized behavior in quantum algorithms. All other operations of a quantum computer are completely deterministic.

§2 A Single Quantum Bit

The operations of a quantum computer allow reading, writing, or manipulating the content of the memory, and therefore serve the same purpose as the operations of a classical computer. The main distinction is that the operations of a quantum computer are formulated to be conformant with the laws of quantum mechanics. We explain in this section the basic operations on a single quantum bit, and introduce some convenient notations.

The input operation of a quantum computer can prepare the memory in any basis state. As a result, each quantum bit is either in the state $|0\rangle$ or in the state $|1\rangle$, but not in a superposition of these basis states. The actual computation is done by applying simple operations, called **quantum gates**, which allow to manipulate the content of the memory. The result of the computation is determined by **measurement operations**.

If the memory consists of a single quantum bit, then the operations are particularly easy to understand. We recall some mathematical vocabulary to ease our discussion. If $x = (x_{m-1}, \ldots, x_0)^t$ and $y = (y_{m-1}, \ldots, y_0)^t$ are vectors in $\mathbf{C}^m$, then

$$\langle x|y\rangle = \overline{x}_{m-1}y_{m-1} + \cdots + \overline{x}_0 y_0$$

defines a **hermitian product**. We follow the convention that hermitian products are anti-linear in the first argument, and linear in the second.

Exercise 2.5 *Show that the hermitian product is positive definite, that is,* $\langle x|x\rangle \geq 0$ *for all* $x \in \mathbf{C}^m$, *and* $\langle x|x\rangle > 0$ *if* $x \neq 0$.

If $x \in \mathbf{C}^m$, then the **norm** of x is defined by $\|x\| = \sqrt{\langle x|x\rangle}$. A vector x with norm $\|x\| = 1$ is called a **unit vector**. Let $U\colon \mathbf{C}^m \to \mathbf{C}^m$ be a linear map. If $\langle\, Ux|Uy\,\rangle = \langle x|y\rangle$ holds for all $x, y \in \mathbf{C}^m$, then U is called **unitary**.

Exercise 2.6 *Show that a complex* $m \times m$ *matrix* U *is unitary if and only if* $U^{-1} = \overline{U}^t$; *that is, the inverse of a unitary matrix is obtained by transposing the matrix and conjugating the matrix entries.*

Exercise 2.7 *A quantum state is a unit vector. Show that if a linear map* M *maps each unit vector* $x \in \mathbf{C}^m$ *to a unit vector* Mx, *then* M *has to be unitary. This property explains the relevance of unitary maps in quantum computing.*

Exercise 2.8 *Let* $\{u_0, \ldots, u_{m-1}\}$ *and* $\{v_0, \ldots, v_{m-1}\}$ *be orthonormal bases of* $\mathbf{C}^m$. *Let* U *be a linear map such that* $v_i = Uu_i$ *for* $i = 0, \ldots, m-1$. *Show that* U *is unitary.*

We have now the terminology to describe the operations on a single quantum bit. A **quantum gate** changes the state of a single qubit by applying an arbitary unitary map U. We use the following graphical notation for such a quantum gate:

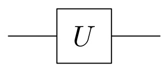

The horizontal line represents the evolution of the quantum bit over time. The time flow is from left to right. The box represents a quantum gate, which applies a unitary map U to the state of the qubit.

The quantum gate is unitary, hence, in particular, linear. This means that the action of the gate is completely determined by its behavior on the base states $|0\rangle$ and $|1\rangle$. Suppose that the quantum gate U changes the input state $|0\rangle$ to $m_{00}|0\rangle + m_{10}|1\rangle$ and the input $|1\rangle$ to $m_{01}|0\rangle + m_{11}|1\rangle$. If the input is a linear combination $a|0\rangle + b|1\rangle$, then the gate U will change this state to

$$\begin{aligned} &a\,(m_{00}|0\rangle + m_{10}|1\rangle) + b\,(m_{01}|0\rangle + m_{11}|1\rangle) \\ &= (am_{00} + bm_{01})|0\rangle + (am_{10} + bm_{11})|1\rangle. \end{aligned}$$

The result of this computation can be expressed in the standard basis (2.1) by the following matrix vector product:

$$\begin{pmatrix} a\,m_{00} + b\,m_{01} \\ a\,m_{10} + b\,m_{11} \end{pmatrix} = \begin{pmatrix} m_{00} & m_{01} \\ m_{10} & m_{11} \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix}.$$

The most familiar example is given by a **not gate**, which changes $|0\rangle$ to $|1\rangle$ and vice versa. This quantum gate can be described by the unitary matrix

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

If we apply this quantum gate twice, then we recover the input. Graphically, we obtain the rule

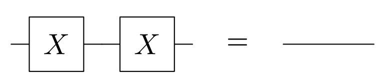

Another operation on one quantum bit is given by the **Hadamard gate**

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

This operation has the following effect:

$$H|0\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle \quad \text{and} \quad H|1\rangle = \frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle.$$

Exercise 2.9 *Calculate $H(H|0\rangle)$ and $H(H|1\rangle)$ by evaluating the expressions in parentheses. Use linearity to obtain the result. Compare your result to the matrix H^2.*

The product of two unitary matrices is a unitary matrix. Therefore, instead of applying gate A and then gate B, we can apply a single quantum gate BA. This way we obtain the rule

$$-\boxed{A}-\boxed{B}- \quad = \quad -\boxed{BA}-$$

The order of the matrices changes because the time flow in a quantum circuit is from left to right. However, the matrices act on column vectors; hence, applying BA means that A is applied first.

Exercise 2.10 *Simplify the circuit, and determine a single unitary matrix Z that is the result of applying the Hadamard gate H, then the not gate X, then again the Hadamard gate H:*

$$-\boxed{H}-\boxed{X}-\boxed{H}- \quad = \quad -\boxed{Z}-$$

Exercise 2.11 *Find a unitary 2×2 matrix R such that*

$$-\boxed{R}-\boxed{R}- \quad = \quad -\boxed{X}-$$

In other words, R should satisfy $R^2 = X$.

Numerous other unitary matrices are used in quantum algorithms. The rotation matrices

$$R(\theta) = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix},$$

and the Pauli matrices σ_x, σ_y, and σ_z are popular choices:

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Exercise 2.12 *Show that the product of any two Pauli matrices is – up to a multiplication by a scalar – either a Pauli matrix or the identity matrix. Memorize the definition of the Pauli matrices.*

An output is obtained by measuring the state of the quantum bit. The **measurement operation** of a quantum bit in the state $a|0\rangle + b|1\rangle$ yields output 0 with probability $|a|^2$, and output 1 with probability $|b|^2$. The state is, in general, changed by the measurement operation. If 0 is observed, then the state is set to $|0\rangle$, and if 1 is observed, then the state is set to $|1\rangle$. We depict a measurement of the quantum bit by a meter sign:

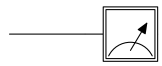

The operations obtained so far allow us to derive a quantum circuit simulating an unbiased coin flip. This circuit produces output 0 with probability 1/2, and output 1 with probability 1/2. We initialize the quantum bit with the state $|0\rangle$, then apply the Hadamard gate, and measure the result:

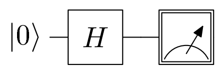

The Hadamard gate changes the state to $\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$; hence, the measurement produces the output with the desired probability.

Exercise 2.13 *Design a quantum circuit that simulates a biased coin flip. The circuit should produce output 0 with probability 1/3, and output 1 with probability 2/3.*

§3 Quantum Gates

We need operations that enable the interaction between different quantum bits. The **xor gate** or **controlled-not gate** acts on two distinct quantum bits. Suppose that the memory contains two quantum bits, then the controlled-not gate operates on the basis states of the system as follows:

$$\begin{array}{rcl} |00\rangle & \mapsto & |00\rangle, \\ |01\rangle & \mapsto & |01\rangle, \\ |10\rangle & \mapsto & |11\rangle, \\ |11\rangle & \mapsto & |10\rangle. \end{array}$$

If we extend this operation linearly, then the quantum state

$$a_{00}|00\rangle + a_{01}|01\rangle + a_{10}|10\rangle + a_{11}|11\rangle$$

will be mapped by this controlled-not gate to

$$a_{00}|00\rangle + a_{01}|01\rangle + a_{10}|11\rangle + a_{11}|10\rangle.$$

Exercise 2.14 *The xor gate is a unitary map. Determine the associated unitary matrix with respect to the computational basis* $|00\rangle, |01\rangle, |10\rangle, |11\rangle$. *Choose the basis vectors in this order.*

Exercise 2.15 *Suppose that a controlled-not gate is applied to the state* $\frac{1}{2}|00\rangle + \frac{1}{2}|10\rangle + \frac{1}{\sqrt{2}}|11\rangle$. *What is the resulting state?*

Controlled-not gates can be generalized to an arbitrary number $n \geq 2$ of quantum bits. A **controlled-not gate** with control bit at position i and target bit at position $j \neq i$ is a unitary map, which is determined by

$$|x_{n-1} \cdots x_1 x_0\rangle \mapsto |y_{n-1} \cdots y_1 y_0\rangle,$$

where x_k and y_k are elements of $\{0, 1\}$, such that $x_k = y_k$ for all $k \neq j$, and the target bit $y_j = x_i \oplus x_j$ is the result of adding x_i to x_j modulo 2. We denote this controlled-not gate by $\Lambda_{i,j}(X)$.

A controlled-not gate $\Lambda_{1,0}(X)$ acting on two quantum bits is depicted in the graphical notation for quantum gates by

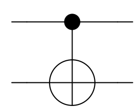

The two horizontal lines represent the two quantum bits. The most significant bit (the bit at position 1) is shown on top, and the least significant bit (the bit at position 0) is shown at the bottom. The black dot $\bullet$ depicts the control bit of the quantum gate, and the crossed circle $\oplus$ depicts the target bit.

Assume that we have three quantum bits, which are initially in the state

$$\frac{1}{2}|001\rangle + \frac{1}{\sqrt{2}}|110\rangle + \frac{1}{2}|111\rangle.$$

Suppose that this state is processed by the quantum circuit

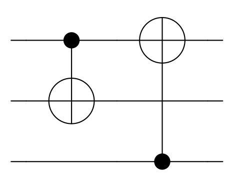

The time flow is from left to right. The first controlled-not gate $\Lambda_{2,1}(X)$ negates the quantum bit in the middle, if the most significant bit is set. The resulting intermediate state after applying the first controlled-not gate is

$$\frac{1}{2}|001\rangle + \frac{1}{\sqrt{2}}|100\rangle + \frac{1}{2}|101\rangle.$$

The second controlled-not gate $\Lambda_{0,2}(X)$ is controlled by the least significant bit, and the target bit is the most significant bit. The intermediate state is changed by this controlled-not gate to

$$\frac{1}{2}|101\rangle + \frac{1}{\sqrt{2}}|100\rangle + \frac{1}{2}|001\rangle.$$

Exercise 2.16 *Design a quantum circuit consisting of controlled-not gates, which realizes the unitary map*

$$|00\rangle \mapsto |00\rangle, \quad |01\rangle \mapsto |10\rangle, \quad |10\rangle \mapsto |01\rangle, \quad |11\rangle \mapsto |11\rangle.$$

We defer the discussion of further multi-qubit operation, and focus instead on operations which act locally on a single quantum bit. It turns out that single-qubit operations and controlled-not gates allow to fully program a quantum computer. Therefore, all other operations can be expressed in terms of these elementary operations. We make a digression and explain tensor products, which provide the proper framework to understand the data structure of the memory.

Let V and W be finite-dimensional complex vector spaces. The tensor product $V \otimes W$ is a vector space, which is spanned by linear combinations of elements $v \otimes w$ such that $v \in V$ and $w \in W$. The product $v \otimes w$ is defined such that it satisfies the additive relations

$$\begin{aligned} (v_1 + v_2) \otimes w &= v_1 \otimes w + v_2 \otimes w && (2.2) \\ v \otimes (w_1 + w_2) &= v \otimes w_1 + v \otimes w_2 && (2.3) \end{aligned}$$

and the balancing relations

$$c(v \otimes w) = (cv) \otimes w = v \otimes (cw) \qquad (2.4)$$

for each v, v_1, v_2 in V, each w, w_1, w_2 in W, and each complex number c.

We can formally construct this vector space $V \otimes W$ as follows. Form the vector space A of all linear combinations of elements (v, w) with $v \in V$ and $w \in W$. Consider the subspace B of A, which consists of all linear combinations of the elements

$$\begin{aligned} (v_1 + v_2, w) - (v_1, w) - (v_2, w), \\ (v, w_1 + w_2) - (v, w_1) - (v, w_2), \\ c(v, w) - (cv, w), \quad c(v, w) - (v, cw), \end{aligned}$$

for $v, v_1, v_2 \in V$, $w, w_1, w_2 \in W$, and $c \in \mathbf{C}$. We define the tensor product $V \otimes W$ to be the quotient space A/B. The image of the element (v, w) of A in $V \otimes W$ is denoted by $v \otimes w$.

We emphasize that not every element of $V \otimes W$ is of the form $v \otimes w$ for some $v \in V$ and $w \in W$. However, every element of $V \otimes W$ can be expressed as a sum $\sum_{i,j} v_i \otimes w_j$ of such tensor products, with $v_i \in V$ and $w_j \in W$.

Exercise 2.17 *Give an example of a vector in $\mathbf{C}^2 \otimes \mathbf{C}^2$ that cannot be written in the form $v \otimes w$ with $v, w \in \mathbf{C}^2$. Prove your result.*

It might be helpful to re-iterate the construction. We started with two finite-dimensional vector spaces V and W. We constructed a giant vector space A with basis $\{(v,w) \mid v \in V, w \in W\}$. The generators of B were chosen such that the quotient space $V \otimes W = A/B$ satisfies the relations (2.2)–(2.4). It is easy to see that $V \otimes W = A/B$ is a finite-dimensional vector space, even though A and B are infinite-dimensional.

Exercise 2.18 *Let V and W be complex finite-dimensional vector spaces. Let $\{e_1, \ldots, e_m\}$ be a basis of V and $\{f_1, \ldots, f_n\}$ be a basis of W. Show that $\{e_i \otimes f_j \mid 0 \leq i < m, 0 \leq j < n\}$ generates $V \otimes W$.*

The exercise shows that $\dim(V \otimes W) \leq \dim(V)\dim(W)$. In fact, it is possible to show that equality holds, which proves that the generating set in the previous exercise is a basis of $V \otimes W$.

Let V and W be as in Exercise 2.18. Suppose that A is a linear map on V, and B is a linear map on W. Let $A \otimes B$ denote the linear map on $V \otimes W$, which is determined by

$$(A \otimes B)(e_i \otimes f_j) = Ae_i \otimes Bf_j.$$

This uniquely determines the values of $A \otimes B$ on other elements of $V \otimes W$ because the elements $e_i \otimes f_j$ are a basis.

Exercise 2.19 *Let A and B be the matrices*

$$A = \begin{pmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \end{pmatrix}, \qquad B = \begin{pmatrix} b_{00} & b_{01} \\ b_{10} & b_{11} \end{pmatrix}$$

representing linear maps with respect to the basis $\{e_0, e_1\}$. Determine the matrix $A \otimes B$ with respect to the basis $\{e_0 \otimes e_0, e_0 \otimes e_1, e_1 \otimes e_0, e_1 \otimes e_1\}$.

The tensor product plays a significant role in quantum computing. Recall that the state space of a single quantum bit is given by $\mathbf{C}^2$. In quantum mechanics, the state space of a joint quantum system is described by the tensor product of the state spaces of its parts. Consequently, a compound system of n quantum bits has the state space

$$\mathbf{C}^2 \otimes \cdots \otimes \mathbf{C}^2 \qquad (n \text{ factors}).$$

This is a 2^n-dimensional complex vector space, hence isomorphic to $\mathbf{C}^{2^n}$. The isomorphism is explicitly given by the linear map

$$|x_{n-1}\rangle \otimes \cdots \otimes |x_1\rangle \otimes |x_0\rangle \longmapsto |x_{n-1}\cdots x_1 x_0\rangle,$$

where $x_i \in \{0,1\}$, $0 \leq i < n$. We will use this isomorphism freely, and switch from one representation to the other, whichever is more convenient. We will silently identify the two notations and write $|00\rangle = |0\rangle \otimes |0\rangle$, etc.

Exercise 2.20 *By convention, the basis vectors associated with the basis* $|0\rangle$ *and* $|1\rangle$ *of* $\mathbf{C}^2$ *are*

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

Derive the vectors of $\mathbf{C}^4 \cong \mathbf{C}^2 \otimes \mathbf{C}^2$ *associated with*

$$|00\rangle = |0\rangle \otimes |0\rangle, \quad |01\rangle = |0\rangle \otimes |1\rangle, \quad |10\rangle = |1\rangle \otimes |0\rangle, \quad |11\rangle = |1\rangle \otimes |1\rangle.$$

Exercise 2.21 *Which vector is associated with* $(a_0|0\rangle + a_1|1\rangle) \otimes (b_0|0\rangle + b_1|1\rangle)$, *assuming the above convention for the basis vectors?*

Suppose we have a memory with n quantum bits. Let U be a unitary 2×2 matrix. We define a **single-qubit gate** U acting on the quantum bit at position i to be the unitary map $\mathbf{1}_{2^{n-i-1}} \otimes U \otimes \mathbf{1}_{2^i}$. Alternatively, one can describe the action of the gate by

$$|x_{n-1}\rangle \otimes \cdots \otimes |x_i\rangle \otimes \cdots \otimes |x_0\rangle \mapsto |x_{n-1}\rangle \otimes \cdots \otimes U|x_i\rangle \otimes \cdots \otimes |x_0\rangle,$$

where $x_i \in \{0,1\}$, $0 \leq i < n$. All tensor components remain unchanged with the exception of $|x_i\rangle$, which is replaced by $U|x_i\rangle$.

Let us illustate this definition in the case of two quantum bits. Suppose that we apply the Hadamard gate H on the least significant bit, that is, the gate acts on the quantum bit at position $i = 0$. The unitary map associated with this gate is represented by the matrix

$$\mathbf{1}_2 \otimes H \otimes \mathbf{1}_1 = \mathbf{1}_2 \otimes H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & -1 \end{pmatrix}. \qquad (2.5)$$

This matrix is the tensor product of the identity matrix $\mathbf{1}_2$ and the Hadamard matrix H.

The alternative description is even easier to grasp. Indeed, the state $|00\rangle = |0\rangle \otimes |0\rangle$ is mapped to

$$|0\rangle \otimes H|0\rangle = |0\rangle \otimes (\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle) = \frac{1}{\sqrt{2}}|0\rangle \otimes |0\rangle + \frac{1}{\sqrt{2}}|0\rangle \otimes |1\rangle.$$

Note that this vector corresponds to the first column of the matrix (2.5). The state $|01\rangle = |0\rangle \otimes |1\rangle$ is mapped to

$$|0\rangle \otimes H|1\rangle = |0\rangle \otimes (\frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle) = \frac{1}{\sqrt{2}}|0\rangle \otimes |0\rangle - \frac{1}{\sqrt{2}}|0\rangle \otimes |1\rangle,$$

and corresponds to the second column of the matrix (2.5). The result of the input $|10\rangle$ and $|11\rangle$ is obtained in a similar way, and we leave these two cases to the reader.

Exercise 2.22 *Suppose that the memory consists of two qubits. Determine the matrix corresponding to the Hadamard gate acting on the most significant qubit.*

The graphical notation for single-qubit gates is similar to the single quantum bit case. A single-qubit gate U acting on the least significant bit in a system of two quantum bits is depicted by

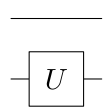

Exercise 2.23 *Determine the action of the circuit*

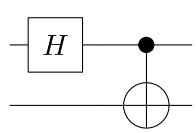

on the input $|00\rangle, |01\rangle, |10\rangle, |11\rangle$. *Explain why the resulting states form an orthonormal basis.*

§4 Measurements

We need to have a way to obtain the value of a quantum bit. The quantum circuit model allows measuring an individual quantum bit with respect to the computational basis. We define these measurement operations in this section and discuss some possible extensions.

Assume that we have a memory consisting of n quantum bits. Suppose that the memory is in the quantum state

$$v = \sum_{x \in \mathbf{F}_2^n} a_x |x\rangle, \qquad a_x \in \mathbf{C}.$$

The state vector v is, as always, assumed to be of unit norm, $\|v\| = 1$. A **measurement** of the quantum bit at position i yields the result $k \in \{0, 1\}$ with probability

$$\sum_{x \in \mathbf{F}_2^n \text{ with } x_i = k} |a_x|^2.$$

The measurement changes, in general, the state vector. If k is observed, then the resulting state of the memory is given by $\frac{1}{\|v_k\|} v_k$, where

$$v_k = \sum_{x \in \mathbf{F}_2^n \text{ with } x_i = k} a_x |x\rangle.$$

Let us illustrate the effect of this operation in the case of two quantum bits. Suppose that the memory is in the state

$$v = \frac{1}{2}|00\rangle + \frac{1}{\sqrt{2}}|10\rangle + \frac{1}{2}|11\rangle.$$

If we measure the qubit at position $i = 0$, then we will observe 0 with probability $(1/2)^2 + (1/\sqrt{2})^2 = 3/4$, and 1 with probability $(1/2)^2 = 1/4$. Note that $v_0 = \frac{1}{2}|00\rangle + \frac{1}{\sqrt{2}}|10\rangle$ and $v_1 = \frac{1}{2}|11\rangle$. Therefore, if we observe 0, then the memory will be in the state

$$\frac{1}{\|v_0\|} v_0 = \frac{2}{\sqrt{3}} v_0 = \frac{1}{\sqrt{3}}|00\rangle + \frac{2}{\sqrt{6}}|10\rangle,$$

and if we observe 1, then the memory will be in the state

$$\frac{1}{\|v_1\|} v_1 = 2v_1 = |11\rangle.$$

Exercise 2.24 *Let* $v = \frac{1}{3}|00\rangle + \frac{\sqrt{3}}{3}|01\rangle + \frac{\sqrt{5}}{3}|10\rangle$. *If we measure the least significant bit, what is the probability to observe 0, respectively 1? Determine the resulting states* v_0 *and* v_1 *of the memory.*

The graphical notation for a measurement is the meter sign. For instance, the measurement of the least significant quantum bit is depicted by

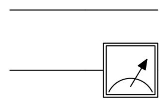

The reader familiar with quantum mechanics will notice that many more types of measurements are, in principle, possible. However, the practical ways to measure quantum bits are typically rather limited. Although quantum physics allows us to measure with respect to any orthonormal basis, we limit ourselves here to the computational basis. If we could perform a measurement with respect to a totally arbitrary orthonormal basis, then there would be no need for quantum gates. The quantum algorithm would then simply consist of a measurement in the appropriate basis.

§5 Examples

We give in this section some tiny examples, which illustrate the notions that we have introduced so far. We will mainly discuss some small quantum circuits, which do not necessarily have any purpose other than illustrating the effect of quantum operations. The superficial examples given here allow us, nonetheless, to illustrate some common tricks of the trade. We will discuss some more meaningful examples in the next chapter.

Example 1. The first example illustrates how the Hadamard gates can be used to generate quickly a superposition of all possible input states. Suppose that the Hadamard gate is applied to both quantum bits, first on the least significant bit, then on the most significant bit:

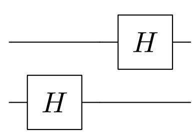

Therefore, the action on the state vector is given by $(H\otimes\mathbf{1}_2)(\mathbf{1}_2\otimes H)$. Suppose that the input is $|00\rangle = |0\rangle\otimes|0\rangle$. The intermediate state after applying the first gate is

$$|0\rangle\otimes H|0\rangle = |0\rangle\otimes(\frac{1}{\sqrt{2}}|0\rangle+\frac{1}{\sqrt{2}}|1\rangle).$$

The final state after applying the second gate is

$$H|0\rangle\otimes(\frac{1}{\sqrt{2}}|0\rangle+\frac{1}{\sqrt{2}}|1\rangle) = (\frac{1}{\sqrt{2}}|0\rangle+\frac{1}{\sqrt{2}}|1\rangle)\otimes(\frac{1}{\sqrt{2}}|0\rangle+\frac{1}{\sqrt{2}}|1\rangle).$$

We can expand the right hand side using the bilinear relations of the tensor product, and obtain the simpler form

$$\frac{1}{2}|00\rangle + \frac{1}{2}|01\rangle + \frac{1}{2}|10\rangle + \frac{1}{2}|11\rangle.$$

We could have obtained the same result by applying the gate on the most significant qubit first, and then the gate on the least significant bit; or even by applying both gates at the same time.

Exercise 2.25 *Suppose that A_1 and B_1 are $n \times n$ matrices, and A_2 and B_2 are $m \times m$ matrices. Show that $(A_1 \otimes A_2)(B_1 \otimes B_2) = (A_1B_1) \otimes (A_2B_2)$.*

A consequence of this exercise is that if we have two quantum gates, which affect disjoint sets of quantum bits, then we can execute these gates in arbitary order. Indeed, we have $(A \otimes \mathbf{1}_m)(\mathbf{1}_n \otimes B) = (\mathbf{1}_n \otimes B)(A \otimes \mathbf{1}_m)$. We can even execute these operations in parallel, because $(A \otimes \mathbf{1}_m)(\mathbf{1}_n \otimes B) = A \otimes B$. Therefore, gates acting on different quantum bits are often denoted on top of each other, as shown on the right, to make the graphical notation more compact:

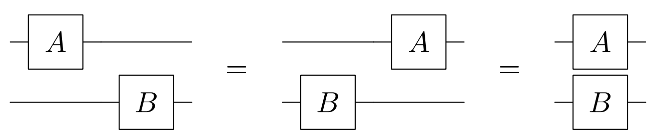

These rules are also useful when one attempts to simplify quantum circuits.

Example 2. Engineering a specific quantum state is a frequent subtask in the design of quantum algorithms. For instance, suppose that we need to prepare four quantum bits in the state

$$\frac{1}{\sqrt{2}}|0000\rangle + \frac{1}{\sqrt{2}}|1111\rangle.$$

Assume that the quantum bits are initially in the state $|0000\rangle$. We can apply the Hadamard gate on the most significant qubit to obtain the state

$$\frac{1}{\sqrt{2}}|0000\rangle + \frac{1}{\sqrt{2}}|1000\rangle.$$

Applying controlled-not gates on the three least significant qubits as target qubits, with the most significant bit as a control bit, yields the desired state

$$\frac{1}{\sqrt{2}}|0000\rangle + \frac{1}{\sqrt{2}}|1111\rangle.$$

Indeed, if we apply the three controlled-not gates to the state $|0000\rangle$, then this state remains unchanged, and if we apply the three controlled-not gates to $|1000\rangle$, then we get $|1111\rangle$; the result follows by linearity of the quantum gates. In graphical notation, the quantum circuit is given by

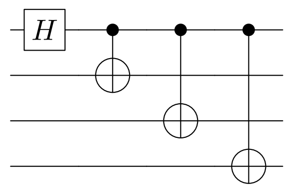

Exercise 2.26 *Design a quantum circuit that prepares the superposition of all basis states with even parity for a system of three quantum bits, namely the state*

$$\frac{1}{2}|000\rangle + \frac{1}{2}|011\rangle + \frac{1}{2}|101\rangle + \frac{1}{2}|110\rangle.$$

Assume that the memory is initially in the state $|000\rangle$*.*

Example 3. Suppose that we have a boolean function $f\colon \mathbf{F}_2^n \to \mathbf{F}_2$. A quantum circuit implementing f has to be realized by a unitary map. This can be accomplished, for instance, by implementing the map

$$|y\rangle \otimes |x\rangle \mapsto |y \oplus f(x)\rangle \otimes |x\rangle$$

on $n+1$ qubits, where $x \in \mathbf{F}_2^n$, and $y \in \mathbf{F}_2$. The most significant bit is the output bit, and the n lowest significant bits are the input bits. The result of $f(x)$ is added modulo 2 to the output bit. The result is a quantum circuit of the form

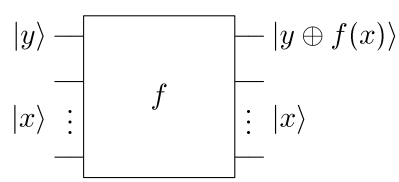

The linearity of the circuit allows to evaluate f for any linear combination of the basis states. Assume that all $n+1$ quantum bits are initialized with state $|0\rangle$. We apply the Hadamard gate to all n input bits. The resulting state is

$$\frac{1}{\sqrt{2^n}} \sum_{x \in \mathbf{F}_2^n} |0\rangle \otimes |x\rangle,$$

a superposition of all possible inputs. If we apply the circuit implementing the function f, then we obtain as a result

$$\frac{1}{\sqrt{2^n}} \sum_{x \in \mathbf{F}_2^n} |f(x)\rangle \otimes |x\rangle.$$

Thus, the circuit evaluates the function f for all possible inputs at once.

Exercise 2.27 *Design a quantum circuit that implements the parity function $f(x_2, x_1, x_0) = x_2 \oplus x_1 \oplus x_0$. Show how this circuit can be used to generate the state $\frac{1}{2}(|0000\rangle + |1010\rangle + |1100\rangle + |0110\rangle)$. Assume that the input is $|0000\rangle$. You can use additional single qubit gates to obtain this result.*

§6 Summary

The state space of a memory with n quantum bits is given by the complex vector space $\mathbf{C}^{2^n} \cong \mathbf{C}^2 \otimes \cdots \otimes \mathbf{C}^2$. We choose, once and for all, a fixed orthonormal basis of this vector space, and call it the computational basis. Its basis vectors are denoted by $|0\cdots 00\rangle, |0\cdots 01\rangle, \cdots, |1\cdots 11\rangle$. An arbitrary state of the memory is of the form

$$\sum_{x \in \mathbf{F}_2^n} a_x |x\rangle, \quad \text{where} \quad \sum_{x \in F_2^n} |a_x|^2 = 1. \tag{2.6}$$

A measurement of the quantum bit at position i yields the result $k \in \{0,1\}$ with probability $\sum_{x_i = k} |a_x|^2$. If k is observed, then the resulting state after the measurement is $\frac{1}{\|v_k\|} v_k$, where v_k denotes the vector

$$v_k = \sum_{x \in \mathbf{F}_2^n, x_i = k} a_x |x\rangle.$$

A single-qubit gate is determined by a matrix $U \in \mathcal{U}(2)$ and a bit position i. Such a gate modifies the state of the memory by applying the unitary matrix $\mathbf{1}_{2^{n-i-1}} \otimes U \otimes \mathbf{1}_{2^i}$. A controlled-not gate $\Lambda_{i,k}(X)$ is specified by its action on the basis vectors

$$\Lambda_{i,k}(X) |x_{n-1} \cdots x_1 x_0\rangle = |y_{n-1} \cdots y_1 y_0\rangle,$$

where $y_j = x_j$ for all $j \neq k$, and $y_k = x_i \oplus x_k$.

Chapter 3

Algorithmic Appetizers

In this chapter, we discuss three small algorithms. The examples illustrate the operations that we introduced in the previous chapter. We begin with a communication protocol, which allows to communicate the state of a single quantum bit. This process is known as teleportation, a somewhat ambitious name for a simple protocol.

§1 Teleportation

Suppose that Alice wants to communicate the state of a quantum bit to Bob. The matter is complicated by the fact that the quantum state might not be known to her. This would not help her much anyway, since, in most cases, she would not be able to communicate a complete description of the state by classical communication alone.

Alice and Bob need, in addition to classical communication, another resource. If Alice and Bob share a pair of quantum bits, which are in the state

$$\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle, \tag{3.1}$$

then it is not difficult to communicate the unknown quantum state, as we will show in this section. This method has been suggested by Bennett, Brassard, Crepeau, Josza, Peres, and Wootters in 1993, and is known as **teleportation**. This type of teleportation has been demonstrated in several experiments.

We need three quantum bits in the teleportation protocol. We assume that the two most significant qubits belong to Alice, and the least significant qubit belongs to Bob. Alice wants to communicate the most significant bit to Bob. We assume that this quantum bit is in the state $a|0\rangle + b|1\rangle$, but Alice might not be aware of that, and the least two qubits are in the state (3.1). Therefore, the system is initially in the state

$$(a|0\rangle + b|1\rangle) \otimes (\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle). \tag{3.2}$$

We assume that Alice and Bob are located far apart. They can apply operations locally on the qubits in their possession and communicate over the phone. The teleportation is surprisingly simple. Alice applies a controlled-not operation $\Lambda_{2,1}(X)$, and a Hadamard gate to the most significant bit. Then she measures her quantum bits, and tells Bob what kind of gate he should apply to his quantum bit.

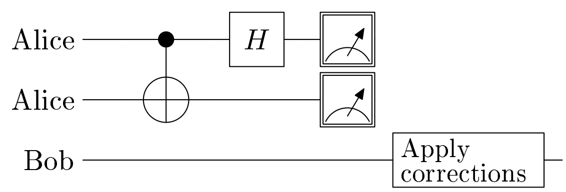

The controlled-not gate $\Lambda_{2,1}(X)$ transforms the state (3.2) to

$$a|0\rangle \otimes (\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle) \;+\; b|1\rangle \otimes (\frac{1}{\sqrt{2}}|10\rangle + \frac{1}{\sqrt{2}}|01\rangle).$$

Applying the Hadamard gate on the most significant qubit yields the state

$$\begin{gathered} a(\tfrac{1}{\sqrt{2}}|0\rangle + \tfrac{1}{\sqrt{2}}|1\rangle) \otimes (\tfrac{1}{\sqrt{2}}|00\rangle + \tfrac{1}{\sqrt{2}}|11\rangle) \\ + b(\tfrac{1}{\sqrt{2}}|0\rangle - \tfrac{1}{\sqrt{2}}|1\rangle) \otimes (\tfrac{1}{\sqrt{2}}|10\rangle + \tfrac{1}{\sqrt{2}}|01\rangle). \end{gathered}$$

The bilinear relations of the tensor product allow this state to be rewritten as follows:

$$\begin{gathered} a(\tfrac{1}{2}|000\rangle + \tfrac{1}{2}|011\rangle + \tfrac{1}{2}|100\rangle + \tfrac{1}{2}|111\rangle) \\ +b(\tfrac{1}{2}|001\rangle + \tfrac{1}{2}|010\rangle - \tfrac{1}{2}|101\rangle - \tfrac{1}{2}|110\rangle). \end{gathered}$$

We collect the terms with the same two most significant qubits, and use the bilinear relations of the tensor product to express this state in yet another, but still equivalent, form:

$$\begin{aligned} \frac{1}{2}\Big(& |00\rangle \otimes (a|0\rangle + b|1\rangle) + |01\rangle \otimes (a|1\rangle + b|0\rangle) \\ & +|10\rangle \otimes (a|0\rangle - b|1\rangle) + |11\rangle \otimes (a|1\rangle - b|0\rangle)\Big). \end{aligned}$$

Alice finally measures the two most significant qubits. The different measurement results and corresponding post-measurement states are shown in the following table:

Observation	Resulting State	Alice tells Bob
00	$\|00\rangle \otimes (a\|0\rangle + b\|1\rangle)$	to do nothing
01	$\|01\rangle \otimes (a\|1\rangle + b\|0\rangle)$	to apply X
10	$\|10\rangle \otimes (a\|0\rangle - b\|1\rangle)$	to apply Z
11	$\|11\rangle \otimes (a\|1\rangle - b\|0\rangle)$	to apply ZX

We note that the resulting state after the measurement can be transformed in each case into a state of the form $|x_2x_1\rangle \otimes (a|0\rangle + b|1\rangle)$, with $x_i \in \{0, 1\}$, by applying the single-qubit gate recommended by Alice. We have accomplished our goal: Alice has successfully communicated the state $a|0\rangle + b|1\rangle$ to Bob.

Entanglement. Let $\mathbf{C}^n$ and $\mathbf{C}^m$ be state spaces of two quantum systems. A state of $\mathbf{C}^n \otimes \mathbf{C}^m$ that can be written in the form $v \otimes w$, for some $v \in \mathbf{C}^n$ and $w \in \mathbf{C}^m$, is called **decomposable**. If a state is not decomposable, then it is called an **entangled state**. Teleportation and many other protocols in quantum computing use entanglement as a resource.

Exercise 3.1 *Show that the state (3.1) is an entangled state.*

There exists a simple criterion that allows us to decide whether an arbitrary state in $\mathbf{C}^2 \otimes \mathbf{C}^2$ is entangled or not. We have to check only a single invariant of the state to decide this question.

Exercise 3.2 *Prove that the state* $|\psi\rangle = \alpha|00\rangle + \beta|01\rangle + \gamma|10\rangle + \delta|11\rangle$ *is decomposable if and only if the coefficients satisfy* $\alpha\delta - \beta\gamma = 0$.

The state (3.1) is called an **Einstein-Podolsky-Rosen state**, or **EPR state** for short. This state received considerable attention after the famous critique on quantum mechanics by Einstein, Podolsky, and Rosen; particularly in Bohm's interpretation. However, there is nothing sacred about this state, and it is, of course, possible to use other entangled states for teleportation.

Exercise 3.3 *Suppose that Alice and Bob share the state* $\frac{1}{\sqrt{2}}|00\rangle + \frac{e^{i\theta}}{\sqrt{2}}|11\rangle$, $\theta \in \mathbf{R}$. *Assume that Alice wants to use her teleport circuit to communicate an unknown state* $a|0\rangle + b|1\rangle$ *of some quantum bit to Bob. Assuming that they both know* θ, *what kind of operations does Bob have to apply when he learns Alice's measurement results? Derive all steps carefully.*

If the state shared by Alice and Bob is not entangled, then teleportation is not possible. However, not every entangled state can be used in for teleportation. We will show later that the shared state has to be a so-called maximally entangled state.

Extensions. Suppose that Alice wants to communicate the state of a system of several quantum bits to Bob. Can she teleport one qubit at a time? We contend that this is the case. To prove this claim, we assume that Alice has $n+1$ quantum bits, which are in the state

$$\sum_{k=0}^{2^n-1}\sum_{j=0}^{1} a_{kj}|k\rangle \otimes |j\rangle \in \mathbf{C}^{2^n} \otimes \mathbf{C}^2. \tag{3.3}$$

If Alice wants to communicate this state to Bob using the teleportation protocol, then she needs to share $n+1$ EPR pairs with Bob. It would be tedious to give a direct proof that this approach works. We show instead that teleportation is faithful in the following sense: If Alice teleports a single qubit, then Alice's remaining n qubits, and the qubit that Bob has received, are in the state (3.3), and these $n+1$ qubits are not entangled with the remaining part of the system. It follows that we can teleport one qubit at a time.

It remains to show that the teleportation of one qubit will preserve the state (3.3), except that one qubit is transferred from Alice to Bob. The intial state of the system is

$$\sum_{k=0}^{2^n-1}\sum_{j=0}^{1} a_{kj}|k\rangle \otimes |j\rangle \otimes (\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle).$$

Note that it suffices to consider one EPR state to teleport a single qubit. We now repeat the exact same teleportation protocol as before. Intially, Alice applies the controlled-not gates $\Lambda_{2,1}(X)$; this yields the state

$$\begin{aligned}\sum_{k=0}^{2^n-1} \Big(& a_{k0}|k\rangle \otimes |0\rangle \otimes (\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle) \\ & + a_{k1}|k\rangle \otimes |1\rangle \otimes (\frac{1}{\sqrt{2}}|10\rangle + \frac{1}{\sqrt{2}}|01\rangle)\Big).\end{aligned}$$

Then she applies the Hadamard gate on the qubit at position 2, which yields the state

$$\begin{aligned}\sum_{k=0}^{2^n-1} \Big(& a_{k0}|k\rangle \otimes \frac{1}{2}(|0\rangle + |1\rangle) \otimes (|00\rangle + |11\rangle) \\ & + a_{k1}|j\rangle \otimes \frac{1}{2}(|0\rangle - |1\rangle) \otimes (|10\rangle + |01\rangle)\Big).\end{aligned}$$

We want to measure the qubits at positions 1 and 2. We use the bilinear relations of the tensor product to rewrite this state in the more convenient, but equivalent, form

$$\begin{aligned}\sum_{k=0}^{2^n-1}\frac{1}{2}\Big(&|k\rangle\otimes|00\rangle\otimes(a_{k0}|0\rangle+a_{k1}|1\rangle)\\ &+|k\rangle\otimes|01\rangle\otimes(a_{k0}|1\rangle+a_{k1}|0\rangle)\\ &+|k\rangle\otimes|10\rangle\otimes(a_{k0}|0\rangle-a_{k1}|1\rangle)\\ &+|k\rangle\otimes|11\rangle\otimes(a_{k0}|1\rangle-a_{k1}|0\rangle)\Big).\end{aligned}$$

Suppose that Alice measures the qubits at positions 2 and 1. If she observes x_2 and x_1, respectively, and informs Bob to apply $Z^{x_2}X^{x_1}$, then after applying Bob's correction operations, we get

$$\sum_{k=0}^{2^n-1}\sum_{j=0}^{1}|k\rangle\otimes|x_2x_1\rangle\otimes a_{kj}|j\rangle=\sum_{k=0}^{2^n-1}\sum_{j=0}^{1}a_{kj}|k\rangle\otimes|x_2x_1\rangle\otimes|j\rangle.$$

We note that Alice's n most significant qubits, and Bob's least significant qubit are in the state (3.3), and that these qubits are not entangled with the qubits at positions 1 and 2.

We can summarize our findings as follows: If Alice wants to communicate the state of $n+1$ quantum bits, then she can do that by applying the teleportation protocol $n+1$ times. If the system is initially in the state

$$\sum_{k=0}^{2^n-1}\sum_{l=0}^{1}a_{kj}|k\rangle\otimes|j\rangle\otimes\bigotimes_{i=0}^{n}(\frac{1}{\sqrt{2}}|00\rangle+\frac{1}{\sqrt{2}}|11\rangle),$$

then after applying $2n+2$ gate operations and $2n+2$ measurements on Alice's side, and up to $2n+2$ operations on Bob's side, they manage to transfer the state (3.3) to Bob.

Remark. Note that the protocol simply communicates quantum states, and it does not teleport matter. You find many exaggerated conclusions in publications about teleportation – watching episodes of Star Trek seems to have side effects.

§2 Deutsch's Problem

Suppose that you are given a black box that contains an implementation of a boolean function $f\colon \mathbf{F}_2 \to \mathbf{F}_2$. Your task is to determine the parity

$f(0) \oplus f(1)$, the sum of $f(0)$ and $f(1)$ modulo 2. The goal is to solve this task with a minimal number of calls to the black box.

The classical solution to this problem requires two calls to the black box, since the function might be constant or not. In the quantum version, you are given an implementation of f as a quantum circuit on two quantum bits,

$$|x_1\rangle \otimes |x_0\rangle \mapsto |x_1\rangle \otimes |x_0 \oplus f(x_1)\rangle, \tag{3.4}$$

with $x_1, x_0 \in \mathbf{F}_2 = \{0, 1\}$. The quantum version can be solved with a single call to the black box. The problem and its solution were suggested by Deutsch in 1985; it is historically one of the first quantum algorithms.

Exercise 3.4 *Give implementations of the quantum circuit (3.4) for the constant functions (a) $f(0) = f(1) = 0$, and (b) $f(0) = f(1) = 1$, as well as for the balanced functions (c) $f(0) = 0$, $f(1) = 1$, and (d) $f(0) = 1$, $f(1) = 0$.*

Let B denote the unitary map on $\mathbf{C}^4$ determined by (3.4). We will derive the solution in some small steps. It is clear that we have to take advantage of the superposition principle to evaluate the boolean function simultaneously for both possible input arguments. The solution to Deutsch's problem uses an additional trick, which allows us to encode the value of $f(x)$ into a phase factor. Suppose that the least significant bit is in the state $1/\sqrt{2}(|0\rangle - |1\rangle)$, then

$$B\left(|x_1\rangle \otimes \big(\frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle\big)\right) = |x_1\rangle \otimes \big(\frac{1}{\sqrt{2}}|f(x_1)\rangle - \frac{1}{\sqrt{2}}|1 \oplus f(x_1)\rangle\big) =: v_{x_1}$$

for all $x_1 \in \{0, 1\}$. If the value of $f(x_1)$ is zero, then the input state remains invariant; otherwise, B affects a change of sign. Explicitly,

$$v_{x_1} = (-1)^{f(x_1)}|x_1\rangle \otimes \big(\frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle\big).$$

We can now use the superposition principle. If we choose $1/\sqrt{2}(|0\rangle + |1\rangle)$ for the most significant qubit, then we obtain the result $1/\sqrt{2}(v_0 + v_1)$ since the black box B is linear. To put this in a different way, we get

$$B\left(\frac{1}{2}(|0\rangle + |1\rangle) \otimes (|0\rangle - |1\rangle)\right) = \frac{1}{2}((-1)^{f(0)}|0\rangle + (-1)^{f(1)}|1\rangle) \otimes (|0\rangle - |1\rangle).$$

The goal was to discriminate between functions, which satisfy $f(0) \oplus f(1) = 0$, and functions satisfying $f(0) \oplus f(1) = 1$. The previous state is equivalent to

$$\begin{cases} \pm\frac{1}{2}(|0\rangle + |1\rangle) \otimes (|0\rangle - |1\rangle) & \text{if } f(0) \oplus f(1) = 0, \\ \pm\frac{1}{2}(|0\rangle - |1\rangle) \otimes (|0\rangle - |1\rangle) & \text{if } f(0) \oplus f(1) = 1. \end{cases}$$

If we apply the Hadamard gate on the most significant qubit, then we get

$$\begin{cases} \pm|0\rangle \otimes \dfrac{1}{\sqrt{2}}(|0\rangle - |1\rangle) & \text{if } f(0) \oplus f(1) = 0, \\ \pm|1\rangle \otimes \dfrac{1}{\sqrt{2}}(|0\rangle - |1\rangle) & \text{if } f(0) \oplus f(1) = 1. \end{cases}$$

We measure the most significant qubit now. If the function in the black box satisfies $f(0) \oplus f(1) = 0$, then we will observe 0 with certainty. If f satisfies $f(0) \oplus f(1) = 1$, then we will observe 1. Note that the algorithm is completely deterministic. We can summarize the algorithm that we have developed as follows:

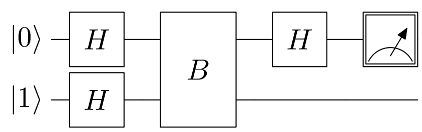

The reader should pause here for a moment and retrace each step in the circuit diagram. The first two Hadamard gates prepare the superposition of the input and the state which allows the encoding of the value of $f(x)$ into a phase factor.

§3 Hidden Subgroup Problems

Deutsch's problem is an instance of a hidden subgroup problem. The hidden subgroup problem is often considered as the Holy Grail of quantum computing and has inspired a considerable amount of research. We need some terminology before we can state this problem. Recall that a **group** is a non-empty set G with a composition operation $\circ\colon G \times G \to G$, such that

G1 $((x \circ y) \circ z) = (x \circ (y \circ z))$ holds for all $x, y, z \in G$;

G2 there exists an element $e \in G$ such that $e \circ x = x \circ e = x$ for all $x \in G$;

G3 for each $x \in G$, there exists an $x^{-1} \in G$ such $x \circ x^{-1} = x^{-1} \circ x = e$.

Axiom G1 states that the composition is associative, and G2 that there exists an identity (or neutral) element. Note that this identity element is uniquely determined. The axiom G3 states that each element x in G has an inverse element.

Exercise 3.5 *Show that (a) the integers* $\mathbf{Z}$ *with addition as composition is a group; (b) the set* $\mathbf{Z}/n\mathbf{Z} = \{0, \cdots n-1\}$ *of integers with addition modulo* n *is a group; (c) the set* $\mathrm{GL}(n, \mathbf{R})$ *of all real invertible* $n \times n$ *matrices is a group with matrix multiplications as composition. Explicitly determine the inverses and the identity element in all cases.*

A subset H of G is called a **subgroup** of G if and only if it forms a group under the restriction of the composition $\circ$ to H. If S is a subset of G, then $\langle S\rangle$ denotes the smallest subgroup of G containing S. If there exists a finite set S such that $\langle S\rangle = G$, then G is called a **finitely generated group**.

Exercise 3.6 *Determine all subgroups of the group* $\mathbf{Z}/6\mathbf{Z}$*.*

Exercise 3.7 *Determine which of the following groups are finitely generated: (a) the additive group of integer* $\mathbf{Z}$*, (b) the group* $\mathbf{Z}/n\mathbf{Z}$*. If possible, give an explicit set of generators.*

We can formulate the problem as follows:

> **The Hidden Subgroup Problem:** Let $f\colon G \to X$ be a black box function from a finitely generated group G to a finite set X such that
>
> $$f(x) = f(y) \quad \text{if and only if} \quad y^{-1}x \in H, \tag{3.5}$$
>
> where H is some initially unknown subgroup of G. Your task is to find a generating set S of H.

The hidden subgroup problem serves as a yardstick measuring the progress in quantum computing. Various instances have been solved, and some see numerous examples in the following chapters.

We have already mentioned that Deutsch's problem can be viewed as a special case of the hidden subgroup problem. Indeed, let the group $G = \mathbf{Z}/2\mathbf{Z}$ and the set $X = \mathbf{Z}/2\mathbf{Z}$. We have two possible subgroups of G, namely $H = \{0, 1\}$ and $H = \{0\}$. If the hidden subgroup $H = \{0\}$, then the constraint (3.5) implies that f has to be a balanced function. If $H = \{0, 1\}$, then f has to be a constant function.

Exercise 3.8 *Assume that* $f\colon \mathbf{Z}/4\mathbf{Z} \to \mathbf{Z}/2\mathbf{Z}$ *is a black box function for a hidden subgroup problem. Enumerate all potential hidden subgroups* H *of* $G = \mathbf{Z}/4\mathbf{Z}$ *that can be encoded by black box functions of this type.*

Exercise 3.9 *Let $f\colon \mathbf{Z}/4\mathbf{Z} \to \mathbf{Z}/2\mathbf{Z}$ be a black box function for a hidden subgroup problem. Assume that the black box is realized by a quantum circuit, which implements the map $|x_1x_0\rangle \otimes |y\rangle \mapsto |x_1x_0\rangle \otimes |y \oplus f(x_1, x_0)\rangle$, with $x_1, x_0, y \in \mathbf{F}_2$. We assume that the binary string x_1x_0 encodes the number $2x_1+x_0$. Design a quantum circuit, which solves this hidden subgroup problem.*

Almost all quantum algorithms that have an exponential speed-up over the best classical algorithms known to date can be formulated as hidden subgroup problems, or some closely related variation of this problem.

§4 A Small Search Algorithm

Suppose that we are given a black box function $f\colon \mathbf{F}_2^n \to \mathbf{F}_2$ such that $f(s) = 1$ for some $s \in \mathbf{F}_2^2$, and $f(x) = 0$ otherwise. We want to find this element s satisfying the search criterion $f(s) = 1$. Classically, we need to evaluate $f(x)$ more than two times to find s with probability greater than $1/2$. We discuss in this section a quantum algorithm that allows us to find s with probability 1 using a single evaluation of the black box function.

We assume that the black box function is given in form of a quantum circuit, which realizes the unitary map B_f given by

$$|x_1x_0\rangle \otimes |y\rangle \mapsto |x_1x_0\rangle \otimes |y \oplus f(x_1, x_0)\rangle,$$

where $x_1, x_0, y \in \mathbf{F}_2$. We evaluate B_f on a superposition of all inputs, and encode the result as a sign change. We accomplish this by initializing with $|0\rangle \otimes |0\rangle \otimes |1\rangle$, and by applying Hadamard gates to all three qubits; these operations generate the state

$$\frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle) \otimes \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle).$$

Applying B_f to this state yields one of the following four possible results:

$f(s) = 1$	resulting state
$s = 00$	$\frac{1}{2}(-\lvert 00\rangle + \lvert 01\rangle + \lvert 10\rangle + \lvert 11\rangle) \otimes \frac{1}{\sqrt{2}}(\lvert 0\rangle - \lvert 1\rangle)$,
$s = 01$	$\frac{1}{2}(\lvert 00\rangle - \lvert 01\rangle + \lvert 10\rangle + \lvert 11\rangle) \otimes \frac{1}{\sqrt{2}}(\lvert 0\rangle - \lvert 1\rangle)$,
$s = 10$	$\frac{1}{2}(\lvert 00\rangle + \lvert 01\rangle - \lvert 10\rangle + \lvert 11\rangle) \otimes \frac{1}{\sqrt{2}}(\lvert 0\rangle - \lvert 1\rangle)$,
$s = 11$	$\frac{1}{2}(\lvert 00\rangle + \lvert 01\rangle + \lvert 10\rangle - \lvert 11\rangle) \otimes \frac{1}{\sqrt{2}}(\lvert 0\rangle - \lvert 1\rangle)$.

We note that the four states are orthogonal. Therefore, it is possible to find a base change T transforming the two most significant qubits into the computational bases states. The coordinate transform is given by

$$T = \frac{1}{2}\begin{pmatrix} -1 & 1 & 1 & 1 \\ 1 & -1 & 1 & 1 \\ 1 & 1 & -1 & 1 \\ 1 & 1 & 1 & -1 \end{pmatrix}.$$

After this base change, we can measure the result in the computational basis. The resulting circuit is

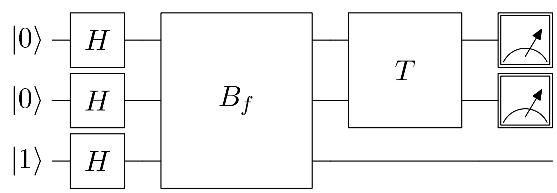

Notice that if the search string is $s = (x_1, x_0)$, then we will observe the two bits (x_1, x_0) in the measurement.

It remains to realize the base change T by a sequence of quantum gates. Note that

$$T = (H \otimes H)\,\mathrm{diag}(1, -1, -1, -1)(H \otimes H).$$

This is easily verified by a direct computation. The diagonal matrix $D = \mathrm{diag}(1, -1, -1, -1)$ can be realized by the circuit

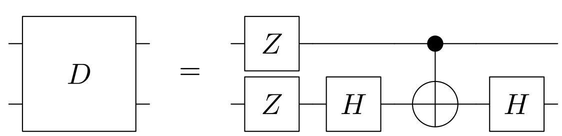

Therefore, we can implement T by

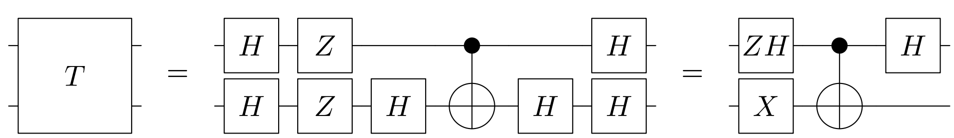

It is possible to generalize this search problem to n quantum bits. A quantum algorithm to solve this problem was published by Grover in 1996. We will discuss his algorithm in detail in one of the following chapters.

§5 Summary

- **Teleportation** is a communication protocol that allows to communicate the state of n quantum bits from Alice to Bob, if they share n EPR pairs.

- **Deutsch's problem** asks to evaluate the parity $f(0) \oplus f(1)$ of a boolean black box function $f\colon \mathbf{F}_2 \to \mathbf{F}_2$. A quantum algorithm can solve this task with a single evaluation of the black box function.

- The **hidden subgroup problem** asks us to find a generating set of an unknown subgroup H of a finitely generated group G, given a black box function f that maps elements of the group G to a finite set X such that $f(x)$ and $f(y)$ are the same if and only if $y^{-1}x \in H$.

- Let $f\colon \mathbf{F}_2^2 \to \mathbf{F}_2$ be a black box function, which is constant zero except on one argument. The **search algorithm** allows us to find this argument with a single evaluation of f. This algorithm was suggested by Grover in 1996.

Bibliography

[1] Sheldon Axler. *Linear Algebra Done Right*. Springer-Verlag, New York, 2nd edition, 2001. (Corrected fourth printing).

[2] Serge Lang. *Linear Algebra*. Springer-Verlag, New York, 3rd edition, 1987.