II. ELEMENTARY TASKS AND TOOLS OF QUANTUM INFORMATION PROCESSING

A. Introduction

Information is quantized in classical digital informations processing as well as in quantum information processing. In analogy to the classical bit, the elementary quantum of information in quantum information processing is called a *qubit*. In the first part of this chapter we will learn how qubits can be implemented by means of spins.
Once the information is stored in a set of qubits, we must be able to manipulate these qubits in order to process the information. This means we must be able to change the state of a qubit either unconditionally (for example, for initializing a qubit or for writing information into a qubit), or conditionally, depending on the previous state of the qubit itself (e.g. the NOT operation) or on the state of another qubit (e.g. the controlled NOT, or CNOT operation), or on the state of the qubit itself and that of another qubit (e.g. the AND operation), and so on. These tasks will have to be performed by *quantum gates*.
Of course one could imagine still more complicated gates, where the state change of one (or more) qubit(s) would depend on the state(s) of an arbitrary number of other qubits. Fortunately all possible operations can be reduced to a finite set of *universal quantum gates*. From these gates one can construct the specific algorithms of quantum information processing which we will study later in this course.

B. Qubits and Spins 1/2

Quantum systems have many possible states all of which obey the principle of superposition. The set of all states of an isolated quantum system thus forms a linear vector space over the complex numbers, the *Hilbert space*. (Note that we are discussing exclusively pure states right now. Mixed states will be introduced later as necessary.) The dimension of that space is related (but not necessarily equal) to the number of values that the energy of the system can assume. Remember that even simple systems may have infinitely many energy eigenstates. The hydrogen atom, for example, has countably infinitely many discrete energy eigenvalues (which are degenerate to different degrees, as discussed in standard quantum mechanics courses), plus uncountably many energy eigenvalues in the continuous spectrum.
In this course we will usually not deal with the continuous spectrum at all, because we want to keep the mathematics simple. The continuous spectrum may be unavoidable when we treat the important subject of decoherence later on. In fact we want to keep the discussion still simpler and we will treat a system which is even simpler than hydrogen without the continuum, namely, a single spin-1/2 particle (with a magnetic moment) fixed in space and subject to an external magnetic field. It is well known that in this case the magnetic moment can only have two different orientations with respect to the external field, which are usually called "up" and "down".
The Hilbert space of a single spin-1/2 particle is thus spanned by the states "up"

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} = |\uparrow\rangle = |0\rangle$$

and "down"

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

(The identification with $|0\rangle$ and $|1\rangle$ follows the convention of [3].)

C. Single-qubit gates

All operators in the Hilbert space of a single qubit can be combined from the four fundamental operators

$$\mathbf{P}_\uparrow = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = |\uparrow\rangle\langle\uparrow|$$

$$\mathbf{P}_\downarrow = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = |\downarrow\rangle\langle\downarrow|$$

$$\mathbf{S}^+ = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} = |\uparrow\rangle\langle\downarrow|$$

$$\mathbf{S}^- = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} = |\downarrow\rangle\langle\uparrow|.$$

More convenient for the purposes of physics are the following combinations:

$$\mathbf{1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \mathbf{P}_\uparrow + \mathbf{P}_\downarrow$$

$$\mathbf{S_z} = \frac{1}{2}\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = \frac{1}{2}(\mathbf{P}_\uparrow - \mathbf{P}_\downarrow) = \frac{1}{2}\mathbf{Z}$$

$$\mathbf{S_x} = \frac{1}{2}\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = \frac{1}{2}(\mathbf{S}^+ + \mathbf{S}^-) = \frac{1}{2}\mathbf{X}$$

$$\mathbf{S_y} = \frac{1}{2}\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} = \frac{i}{2}(\mathbf{S}^- - \mathbf{S}^+) = \frac{1}{2}\mathbf{Y}$$

$$\mathbf{S_x}^2 = \mathbf{S_y}^2 = \mathbf{S_z}^2 = \frac{1}{4}\mathbf{1}.$$

These "spin matrices" can be used to write the Hamiltonian of a spin-1/2 particle (fixed in space) in an external field with components B_x, B_y, B_z:

$$\mathbf{H} = -\vec{B}\cdot\vec{\mathbf{S}} = -(B_x\mathbf{S_x} + B_y\mathbf{S_y} + B_z\mathbf{S_z}).$$

where all g factors etc have been absorbed in the units of $\vec{B}$.

Any unitary 2×2 matrix is a valid single-qubit quantum gate. Note that the operators $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$ have eigenvalues ± 1 and thus are unitary. It is evident why $\mathbf{X}$ is also often called the "NOT gate" in the language of quantum computing; $\mathbf{Z}$ generates a π relative phase between the two basis states, and $\mathbf{Y}$ is a combination of the two other gates. It is also easy to generate an arbitrary relative phase (instead of π) between the two states. To see this, note that

$$\exp(i\phi\mathbf{Z}) = \begin{pmatrix} e^{i\phi} & 0 \\ 0 & e^{-i\phi} \end{pmatrix},$$

which generates a relative phase 2ϕ. The NOT gate can also be generalized. Due to the fact that $\mathbf{X}^2 = \mathbf{1}$ we have

$$\exp(i\phi\mathbf{X}) = \mathbf{1}\cos\phi + i\mathbf{X}\sin\phi = \begin{pmatrix} \cos\phi & i\sin\phi \\ i\sin\phi & \cos\phi \end{pmatrix},$$

which interpolates smoothly between the identity and NOT gates, for $\phi = 0$ and $\frac{\pi}{2}$, respectively. For $\phi = \frac{\pi}{4}$ we obtain a "square root of NOT" gate.

It is plausible (and, using the algebraic properties of the spin matrices, easy to show) that for any unit vector $\hat{n}$, the square of (twice) the spin component operator along $\hat{n}$ is the unit operator,

$$(2\hat{n}\cdot\vec{\mathbf{S}})^2 = (n_x\mathbf{X} + n_y\mathbf{Y} + n_z\mathbf{Z})^2 = \mathbf{1},$$

and consequently

$$\mathbf{R}_{\hat{n}}(\theta) = \exp(i\theta 2\hat{n}\cdot\vec{\mathbf{S}}) = \mathbf{1}\cos\theta + 2\hat{n}\cdot\vec{\mathbf{S}}\sin\theta.$$

This operator obviously commutes with the spin component $\hat{n}\cdot\vec{\mathbf{S}}$ and thus does not affect this specific component. In fact it can be shown that the unitary transformation $\mathbf{R}_{\hat{n}}(\theta)$ corresponds to a rotation (by the angle 2θ, in fact) in spin space about the axis $\hat{n}$. Note that a 2π rotation ($\theta = \pi$) reverses the sign of any single-qubit state, but has no consequences for expectation values of physical observables in that state. Any unitary single-qubit operator can be written in the form

$$\mathbf{U} = e^{i\alpha}\mathbf{R}_{\hat{n}}(\theta).$$

It is often desirable to employ only rotations about the coordinate axes instead of rotations about arbitrary axes $\hat{n}$. This is indeed possible:

$$\mathbf{U} = e^{i\alpha}\mathbf{R}_{\hat{z}}(\beta)\mathbf{R}_{\hat{y}}(\gamma)\mathbf{R}_{\hat{z}}(\delta).$$

For any unitary $\mathbf{U}$ suitable values of $\alpha\cdots\delta$ can be found. A similar decomposition with $\hat{x}$ instead of $\hat{z}$ can also be found. More details on quantum gates (not only for a single qubit) can be found in our lecture notes from last semester and in Chapter 4 of Nielsen and Chuang [3].

End of week 2, April 26.

D. Controlled gates

Any programing language contains control structures of the type: "If condition X holds, perform operation Y". The simplest implementation of that structure is the 2-bit (or 2-qubit) operation known as "controlled not" (CNOT), defined by the following truth table:

control-qubit	target-qubit	result
0	0	00
0	1	01
1	0	11
1	1	10

The control qubit remains unchanged, but the target qubit is flipped if the control qubit is 1. (We abbreviate $|1\rangle$ as 1 here for simplicity.) The "result" column of the truth table lists both control and target qubits. Note that the output target qubit is equal to the "exclusive or" (XOR) between the control and target qubits. Hence the CNOT operation is also called "reversible XOR", where the reversibility is accomplished by keeping the value of the control qubit, in contrast to the ordinary (irreversible) XOR operation of classical computer science. In fact, the reversible XOR is its own inverse. Symbolically it achieves the following mapping:

$$(x, y) \longrightarrow (x, x \text{ XOR } y),$$

and it can be used to copy a bit, because it maps

$$(x, 0) \longrightarrow (x, x).$$

A combination of three CNOT gates (the second one with reversed roles of control and target bits) swaps the contents of two bits. Thus the CNOT gate can be used to copy and move bits around. In matrix notation with respect to the usual computational basis $(|00\rangle, |01\rangle, |10\rangle, |11\rangle)$ the CNOT gate reads

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} = \begin{pmatrix} \mathbf{1} & \mathbf{0} \\ \mathbf{0} & \mathbf{X} \end{pmatrix}$$

(using an obvious block matrix notation). Replacing $\mathbf{X}$ by an arbitrary unitary single-qubit operation $\mathbf{U}$, we arrive at the more general *controlled-U* (CU) gate which can be implemented using CNOT and single-qubit gates.

In *higher-order controlled operations* n control qubits and k target qubits are used; an important example is the Toffoli (controlled-controlled-NOT, or C^2NOT) gate, or more general, the C^2U gate for some arbitrary single-qubit $\mathbf{U}$. This gate performs the $\mathbf{U}$ operation on the single target qubit only if both of the control qubits are set to 1. Actually, C^2U can be built from CNOT and single-qubit gates. (For details, see Chapter 5 of last semester's lecture notes.)

E. Universal quantum gates

It is important to know if any conceivable unitary operation in the Hilbert space of interest can be decomposed into a sequence of standard elementary operations taken from a finite set. Only if that is true, a universal quantum computer can be built which can be programmed to fulfill fairly arbitrary tasks, much as today's universal classical digital computers which are (in principle) built from a very small set of universal classical gates. Luckily there exists a set of *universal quantum gates*, in the sense that any unitary operation may be *approximated* to arbitrary accuracy by a combination of these gates.
The following four gates do the trick

- CNOT
- The $\frac{\pi}{8}$ gate

$$\mathbf{T} = \begin{pmatrix} 1 & 0 \\ 0 & \exp i\frac{\pi}{4} \end{pmatrix} = \exp\Big(i\frac{\pi}{8}(\mathbf{1}-\mathbf{Z})\Big).$$

- The phase gate

$$\mathbf{S} = \mathbf{T}^2 = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}.$$

  (Note that $\mathbf{S}^2 = \mathbf{Z}$.)

- The Hadamard gate

$$\mathbf{H} = \frac{1}{\sqrt{2}}(\mathbf{X}+\mathbf{Z}) = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

This set of four gates can be shown to be universal in a three-step process.
i) Any unitary operator can be expressed (exactly) as a product of unitary operators affecting only two computational basis states: "Two-level gates are universal."
ii) (From i) and preceding subsections.) Any unitary operator may be expressed (exactly) using single-qubit and CNOT gates: "Single-qubit and CNOT gates are universal."
iii) Single-qubit operations may be *approximated* to arbitrary accuracy using Hadamard, phase, and $\frac{\pi}{8}$ gates.
Here we will only give some brief comments on these three steps; more details can be found in Chapter 5 of last semester's lecture notes.
Step i) is fairly simple. Any $d \times d$ unitary matrix $\mathbf{U}$ can be written as a product of (at most) $\frac{d(d-1)}{2}$ two-level unitary matrices (that is, unitary matrices which act non-trivially only on at most two vector components). One can multiply the matrix $\mathbf{U}$ by suitable two-level unitary matrices to eliminate eliminate non-diagonal elements of successively, until $\mathbf{U}$ has been reduced to the unit matrix. The original $\mathbf{U}$ can then be reconstructed from the inverses of the matrices used to eliminate the elements.
Step ii) is also not difficult. Single-qubit and CNOT gates can be used to build the arbitrary two level gates discussed in the previous step. The basic idea is simple: Transform the Hilbert space such that the two relevant basis states become the basis states of one qubit, perform the desired single-qubit operation on that qubit, and transform back to the original basis. The basis reshuffling can be achieved via higher-order controlled-NOT operations, which in turn can be reduced to simple CNOT operations.
Step iii) also is not too difficult to understand in principle, but the actual proof involves some rather tricky mathematics unfamiliar to most physicists. Remember that the operations $\mathbf{R}_{\hat{n}}(\beta)$ had an interpretation as rotations in spin space, or on the so-called *Bloch sphere* which is a way to represent all normalized single-qubit states. It is clear that any point of a sphere can be mapped to any other point by combining two rotations (by two arbitrary angles) about two mutually orthogonal axes. Imagine we could implement a rotation about some axis $\hat{n}$ by an angle α which is an *irrational* multiple of 2π. Due to irrationality, the angles

$$n\alpha \mod 2\pi \quad (n = 0(1)\infty)$$

are dense in $[0, 2\pi]$ and thus an arbitrary rotation about $\hat{n}$ can be approximated to arbitrary precision by repeating the α rotation:

$$\mathbf{R}_{\hat{n}}(\beta) = (\mathbf{R}_{\hat{n}}(\frac{\alpha}{2}))^{\nu} + O(\epsilon).$$

It is not too difficult to show that the Hadamard, $\frac{\pi}{8}$ and phase gates can be used to construct two rotations by the same angle $\alpha = \arccos\left(\cos^2\left(\frac{\pi}{8}\right)\right)$ about two mutually orthogonal axes. Some theorems from algebra and number theory are then used to show that α is indeed an irrational multiple of π.