Introduction to Quantum Information Processing — Grundlagen zur Quanteninformationsverarbeitung

Joachim Stolze and Dieter Suter
Institut für Physik, Universität Dortmund
WS 2001/02

November 15, 2001

Abstract

A two-hour introductory course mainly designed for the graduate students of the "Graduiertenkolleg 726: Materialeigenschaften und Konzepte für die Quanteninformationsverarbeitung". Each section (hopefully) corresponds to one week of the semester.

1 Introduction and survey

A lot of material related to this course will be presented at the introductory workshop on October 16, starting 10.30 in Hörsaal, 2 HG II. Students of this course are supposed to attend the workshop, which replaces the 18 October session of this course.
General references for this course are the books [1, 2, 3, 4] and also the review articles [5, 6]. The book [3] by Nielsen and Chuang claims to be an introductory textbook for both physicists and computer scientists. An excellent, though formal treatment of the theoretical aspects is provided by Preskill's lecture notes [7] available on the web . The web site also contains an updated literature list with commentary. Of course you may also browse the material available at the homepage of our seminar on quantum computing.

1.1 Motivation: why quantum computers

Over the past twenty years, much of the world economy has been driven by the development in microelectronics and information technology.

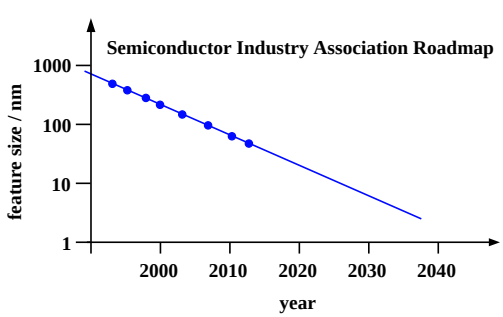

Figure 1.1: Prospective evolution of feature size in microelectronic circuits (source: international semiconductor association roadmap).

The feature size of electronic devices is now in the range of 100 nm and decreasing at a rate of some 12% per year. Acording to this roadmap, feature sizes of 50 nm will be reached in the year 2013. At the same time the operational voltage decreases to less than one Volt.
The capacity of a spherical capacitor is $C = 4\pi\epsilon_0 r$. For a spherical capacitor with radius 50 nm, the capacity is therefore of the order of $5 \cdot 10^{-18} C$. A change in the voltage of 0.1 V will then move less than four electrons

in such a device. While intelligent designs will be able to avoid problems from such effects, the quantisation of charge will certainly become noticeable in electronic devices around the year 2020. By that time, quantum mechanics will have to be considered explicitely when the design and function of electronic devices is analysed. While this may appear as a nuisance to many engineers, it also represents an enormous potential, since devices that incorporate quantum mechanical effects into their principle of operation may be much more powerful than conventional (classical) devices. The emerging resarch field of quantum information processing (qiv) studies the possibilities provided by such an approach and tries to develop algorithms and implementations to harness this potential.

1.2 History of quantum information processing

Quantum mechanics has always been the basis for understanding the properties of semiconductor materials, which form the basis of today's computers. However, if one is not interested in the first principles derivation of these materials properties, it is well possible to describe the operation of todays computers by classical electrodynamics.
It was soon recognized that the situation might be different when the components keep shrinking in size, thus approaching atomic dimensions. In such a situation, quantum mechanics will certainly have to be taken into account explicitely. Quantum effects may prevent current electronic circuits from functioning, e.g. when the Coulomb blockade prevents current from flowing. On a more fundamental level, one may ask how quantum mechanics can describe Boolean logic. In particular, logical gates like AND are clearly not reversible, since it is not possible to calculate the input from the output. Quantum mechanics, on the other hand, describes all time evolution with unitary operators, such as $exp(iHt)$, which are obviously time-reversible ($exp(iHt)exp(-iHt) = 1$).
Several people, such as Bennett, Fredkin, and Toffoli started to discuss the possibility of describing computer operation quantum mechanically. They showed, e.g., that algorithms based on Boolean logic can also be adapted to reversible logic. In 1982, Benioff [8] showed that quantum mechanical systems can efficiently simulate classical systems. This proved that quantum mechanical systems are at least as powerful as classical systems.
In the same year, Richard Feynman asked the opposite question: Can classical computers efficiently simulate quantum mechanical systems [9].

1982 Richard Feynman

R.P. Feynman, *'Simulating physics with computers'*, Int. J. theor. Phys. 21, 467-488 (1982).

Die Rechenleistung, die benötigt wird, um quantenmechanische Systeme zu simulieren, wächst exponentiell mit der Größe des Systems

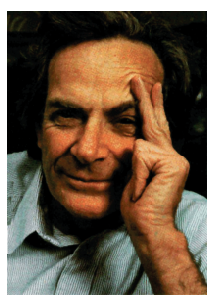

Man benötigt einen Computer, der selber ein Quantensystem ist.

He noted that the number of variables required to describe the system grows exponentially with its size. As an example consider a system of N spins 1/2. The size of the corresponding Hilbert space is 2^N and a specification of its wavefunction therefore requires $2^{2N} - 1$ real numbers. On any given (classical) computer, the duration of a simulation will therefore grow exponentially with the number of particles in the quantum system. He concluded that classical computers will never be able to exactly simulate quantum mechanical systems containing more than just a few particles. Of course, these considerations only take the general case into account. If the particles (or at least the majority) do not interact, e.g., it is always possible to perform the computation in a smaller Hilbert space, thus reducing the computational requirements qualitatively.

After stating the problem, Feynman immediately offers a solution: "Quantum computers - universal quantum simulators". He shows that this drastic increase in the computation time is a direct consequence of the large amount of information in the quantum mechanical system. This implies that quantum systems are efficient processors of information. He states "I therefore believe it is true that with a suitable class of quantum machines you could imitate any quantum system, including the physical world." As an open question he asks which systems could actually be simulated and where it would be useful.
A first proof of this conjecture was given in 1993 by Bernstein, Vazirani, and Yao. They showed that a quantum mechanical Turing machine was capable of simulating other quantum mechanical systems in polynomial time. In the following years, several people developed algorithms that would run more efficiently on quantum computers than on classical machines. However, these examples were relatively artificial and of little practical relevance. The first example of an algorithm that used the power of quantum computers and may have practical relevance was provided by Peter Shor in 1994 [10]. He discusses (amongst other problems) an algorithm for the factorization of prime numbers. The best classical algorithms for factorization of an l digit number use a time that grows as $\exp(cl^{(1/3)}(\log l)^{(2/3)})$, i.e. exponentially with the number of digits. He proposes a model for quantum computation and an algorithm that solves the factorization problem in a time proportional to $O(l^2 \log l \log\log l)$, i.e. polynomially in time. This is a qualitative difference: algorithms that are polynomial in time are considered "efficient", while exponential algorithms are not useable for large systems. The different bahavior will always make classical computers slower than quantum computers at some system size, independent of the speed at which individual operations are performed. This can be seen by putting in some numbers for the above example. We will assume that a fast classical computer can factorize a 50 digit number in one second, while the quantum computer may take as much as an hour for the same operation. Setting the constant c in the above formula to 1 (it was not specified in Shor's paper),
If the number of digits increases to 300, both computers require some 2.5 days to solve the problem. A further increase to 1000 digits requires 42 days on the quantum computer, while the classical computer would need some 19000 years - clearly too long for any practical purposes. With 2000 digits, the quantum computer needs half a year, while the computation time on the classical computer becomes roughly equal to the age of the universe.

1.3 Qauntum communication

One of the most active areas of quantum information processing is quantum communication, i.e. the transfer of information encoded in quantum mechanical degrees of freedom. Semiclassically, a photon can carry a bit: it can be transmitted or not, thus corresponding to a logical 0 or 1. Other encoding schemes include the polarization of the photon, which may be vertical or horizontal. Quantum mechanically, a photon has to be described by a vector in Hilbert space. If we consider polarization, e.g., the relevant part of Hilber space is the spinor $\psi = a\psi_x + b\psi_y$, where ψ_x andψ_y describe linear polarization states in the two direction. The two parameters a and b are both complex numbers. Taking normalization into account, the system is therefore described by three continuous variables. This does not mean, however, that it is possible to store an infinite amount of information in a single photon. To obtain the information content, one has to take the measurement process into account: it is never possible to exactly measure the quantum state of a single photon. A single measurement can only measure one degree of freedom and returns a single bit (particle found or not). A complete measurement of the state of a single photon would thus require repeated measurements, which were possible if one could prepare copies of the actual quantum mechanical state. However, this is prohibited by the "no cloning theorem", which states that no process can duplicate the exact quantum state of a single particle.
While the details of the calculation are rather involved, it is possible to show that a single quantum mechanical two-level system can transfer up to two classical bits of information. This is just one example that proves the by now famous saying "information is physical", which means that information is not just an abstract concept, but must always be related to the physical system carrying the information.
Quantum communication has evolved into a very active field. Besides the fundamental interest, it promises a number of possible applications: taking quantum mechanics into account may improve the information content of communication channels, as discussed above. In addition, it has been shown that communication with individual photons may be made secure, i.e. it is impossible to tap into such a communication without the users of the communication line noticing it. This is a consequence of the no-cloning theorem: While it is conceivable that an evesdropper intercepts a photon, thus detecting that information is being transferred, and re-emitting the photon to the original receiver, he cannot send an exact copy of the original photon. If the communication protocol were to use only the presence or absence of the photon as the information, the evesdropper would be able to use QND (=quantum nondemolition detection) to observe the passage of the photon. Such experimental schemes can measure a given quantum mechanical variable (such as light intensity) without affecting this variable (i.e. chaning the number of photons). They must, however, change the conjugate variable, in this example the phase of the photon, thus allowing the original recipient to detect the presence of the evesdropper.

1.4 Quantum computer: principle of operation

A quantum computer, i.e. a programmable quantum information processing device, encodes the information in the form of a quantum register, consisting of a labeled series of quantum bits or qubits. Each qubit is represented by a quantum mechanical two level system, such as a spin 1/2. While today's quantum registers are limited to 7 qubits, a useful quantum computer will requireseveral hundred to 1000 qubits. Before an actual computation, the quantum register must be initialized into a well defined state, typically the ground state $|0, 0, ...0>$. Quantum gates operate on this state; the sequence of quantum gates required depends on the specific algorithm. The sequence of operations, i.e. the program, may be stored in a classical device associated with the quantum computer, such as a classical computer. These operations can be represented as unitary transformations, typically of the form $\exp(-iHt)$. It has been shown that a general purpose computer can be built on the basis of the following quantum gates:

- single qubit operations, corresponding to arbitrary rotations $R(\theta, \phi)$
- one type of 2-qubit operations, e.g. the "controlled not" or CNOT

This 2-qubit operation corresponds to the following truth table:

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

Any implementation of a quantum computer must therefore provide the possibility of performing one-qubit operations on arbitrary qubits, without affecting other qubits. It must therefore include a means of addressing individual qubits. The 2-qubit operations must also be applied to arbitrary pairs of qubits. It is possible, however, to decompose a 2-qubit operation between any pair into a series of 2-qubit operations between nearest neighbours. Such schemes are often much easier to implement than schemes with interactions between arbitrary pairs. The number of 2qubit operations is larger, but increases only linearly with the number of qubits. The overall process therefore remains efficient.

1.5 Some quantum mechanics

The following material (compare, for example, Chapter 2 of Ref. [7]) should be a mere recapitulation of the quantum mechanics course for most of you:

- Vectors in Hilbert space
- Operators in Hilbert space
- Dynamics: the Schrödinger equation
- The two-dimensional Hilbert space: Spins and single qubits
- Two qubits and the density matrix

1.5.1 Vectors in Hilbert space

Systems with a finite number of particles in a bounded region of space (non-ionized atoms, for example) have a *discrete spectrum* of energy eigenvalues. For the time being we *neglect* the continuous spectrum which may exist in other systems. It will later make its appearance in a cumbersome way anyway ($\rightarrow$ dissipation, decoherence). For mathematical simplicity we will even assume that the dimension d of the Hilbert space is *finite*. $d = 2$ will be the important special case of a single qubit.
The Hilbert space thus is a d-dimensional complex linear space: every linear combination of allowed states (Hilbert space vectors) is an allowed state too; scalar product, norm, etc can be defined as usual:

d-dimensional (column) vector: $|a\rangle = \begin{pmatrix} a_1 \\ a_2 \\ . \\ . \\ a_d \end{pmatrix}$

corresponding (row) vector $\langle a| = (a_1^*, a_2^*, \ldots, a_d^*)$ (*: complex conjugation)
It suffices to consider normalized states $|\psi\rangle$, that is, $\langle\psi|\psi\rangle = 1$. Furthermore the states $|\psi\rangle$ and $e^{i\alpha}|\psi\rangle(\alpha$ real) are physically equivalent: overall phase factors do not matter. However, *relative* phases between components of a state are *extremely* important:$|\phi\rangle + |\psi\rangle$ and $|\phi\rangle + e^{i\alpha}|\psi\rangle$ may have completely different physical properties.

1.5.2 Operators in Hilbert space

Operators map states (*linearly*) to each other; they thus are $d \times d$ complex matrices operating on the d-dimensional Hilbert space.

$$\mathbf{A}|\psi\rangle = |\phi\rangle$$

Observables, that is, measurable quantities correspond to *self-adjoint* or *Hermitean* matrices, that is,

$$\mathbf{A}^\dagger = \mathbf{A}; \quad (\mathbf{A}^\dagger)_{ij} := (\mathbf{A})^*_{ji}$$

Self-adjoint operators have real eigenvalues (measurable quantities!); their eigenstates are pairwise orthogonal (or can be orthogonalized in the case of degeneracy). Thus they form a *basis* in Hilbert space.

$$\mathbf{A}|a_i\rangle = a_i|a_i\rangle \quad (i = 1, \ldots, d)$$
$$\langle a_i|a_j\rangle = 0 \text{ for } i \neq j$$

... and for normalized states we have

$$\langle a_i|a_j\rangle = \delta_{ij} \quad (\text{ Kronecker-}\delta)$$

Eigenstates and eigenvalues characterize an operator completely, because any arbitrary state can be expanded in eigenstates of $\mathbf{A}$. This leads to the *spectral representation* of $\mathbf{A}$. To define that representation we need a further class of operators: *projection operators* or *projectors* for short.

$$\mathbf{P}_i := |a_i\rangle\langle a_i|$$

Action on an arbitrary state $|\psi\rangle$

$$\mathbf{P}_i|\psi\rangle = |a_i\rangle\langle a_i|\psi\rangle = \underbrace{\langle a_i|\psi\rangle}_{\text{a number}} |a_i\rangle$$

$|\langle a_i|\psi\rangle|$is the length of the projection of $|\psi\rangle$ on the unit vector $|a_i\rangle$:

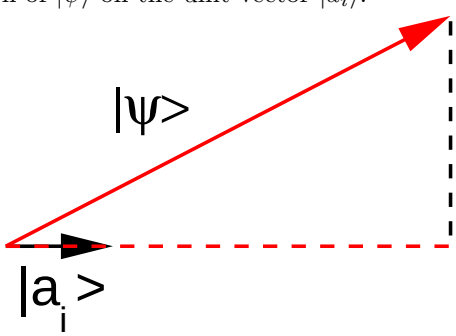

For orthonormal $|a_i\rangle$ (i.e. $\langle a_i|a_j\rangle = \delta_{ij}$) we have

$$\mathbf{P}_i\mathbf{P}_j = \delta_{ij}\mathbf{P}_j; \text{ especially } \mathbf{P}_i^2 = \mathbf{P}_i,$$

and as the $\mathbf{P}_i$ cover "all directions" of Hilbert space:

$$\sum_{i=1}^{d} \mathbf{P}_i = \sum_{i=1}^{d} |a_i\rangle\langle a_i| = \mathbf{1}$$

(completeness). Now we can define the *spectral representation* of $\mathbf{A}$:

$$\mathbf{A} = \sum_{i=1}^{d} a_i\mathbf{P}_i = \sum_{i=1}^{d} a_i|a_i\rangle\langle a_i|$$

An arbitrary state is decomposed into components along eigenstates, each such component is treated accordingly. Postulate: A (single) *measurement* of the observable $\mathbf{A}$ in the (normalized) state $|\psi\rangle$ yields one of the eigenvalues a_i of $\mathbf{A}$ with probability $|\langle a_i|\psi\rangle|^2$ ($\sum_i |\langle a_i|\psi\rangle|^2 = 1$ due to normalization). Immediately after the measurement the system is in the (normalized) state

$$\frac{\mathbf{P}_i|\psi\rangle}{||\mathbf{P}_i|\psi\rangle||}.$$

In general it is not possible to predict the outcome of a measurement. A measurement of $\mathbf{A}$ on an ensemble of systems prepared in the same state $|\psi\rangle$ yields the *average* (expectation value)

$$\langle \mathbf{A} \rangle := \langle \psi | \mathbf{A} | \psi \rangle$$

with deviations described by the *variance*

$$\langle (\mathbf{A} - \langle \mathbf{A} \rangle)^2 \rangle \geq 0 \quad (\text{ "=" for an eigenstate })$$

(Note that we have just discussed two different kinds of measurement.)

1.5.3 Dynamics: The Schrödinger equation

A state $|\psi(t)\rangle$ evolves according to the "time-dependent" *Schrödinger equation*

$$\frac{d}{dt}|\psi(t)\rangle = -\frac{i}{\hbar}\mathbf{H}|\psi(t)\rangle$$

Where $\mathbf{H}$ is the Hamiltonian. If $|\phi_i\rangle$ is an eigenstate of the Hamiltonian with energy eigenvalue ε_i:

$$\mathbf{H}|\phi_i\rangle = \varepsilon_i |\phi_i\rangle$$

("time-independent Schrödinger equation") then

$$|\psi(t)\rangle = \exp\left(-i\frac{\varepsilon_i t}{\hbar}\right)|\phi_i\rangle$$

is a solution of the time-dependent Schrödinger equation with initial condition

$$|\psi(t=0)\rangle = |\phi_i\rangle.$$

Obviously this is a *stationary* state, as a global phase factor has no physical significance. As any initial state $|\psi(t=0)\rangle$ can be expressed as a linear combination of eigenstates $|\phi_i\rangle$ of $\mathbf{H}$, the initial value problem is solved (at least in principle). Formally the solution for **time-independent H** can be written as

$$|\psi(t)\rangle = \exp\left(-i\frac{\mathbf{H}t}{\hbar}\right)|\psi(t=0)\rangle.$$

The *time evolution operator* $\mathbf{U}(t) := \exp\left(-i\frac{\mathbf{H}t}{\hbar}\right)$ may be defined/interpreted in two ways:
i) as a power series

$$\exp\left(-i\frac{\mathbf{H}t}{\hbar}\right) = \mathbf{1} + \left(-i\frac{\mathbf{H}t}{\hbar}\right) + \frac{1}{2}\left(-i\frac{\mathbf{H}t}{\hbar}\right)^2 + \frac{1}{6}\left(-i\frac{\mathbf{H}t}{\hbar}\right)^3 + \cdots$$

ii) by means of the spectral representation

$$\exp\left(-i\frac{\mathbf{H}t}{\hbar}\right) = \sum_{i=1}^{d} \exp\left(-i\frac{\varepsilon_i t}{\hbar}\right)|\phi_i\rangle\langle\phi_i|$$

(NB: For *time-dependent* $\mathbf{H}(t)$ $\mathbf{U}(t)$ is the solution of an operator differential equation; for completely general time dependence the solution of that equation is not even known for $d = 2$.)
All eigenvalues $\exp\left(-i\frac{\varepsilon_i t}{\hbar}\right)$ of $\mathbf{U}(t)$ have unit modulus; operators with this property are called **unitary**. $\mathbf{U}$ preserves all scalar products, that is, the scalar product of $|\psi\rangle$ and $|\chi\rangle$ equals that of $\mathbf{U}|\psi\rangle$ and $\mathbf{U}|\chi\rangle$; consequently *norms* are preserved too. Thus unitary operators are *rotations* in Hilbert space. The general property characterizing unitarity is

$$\mathbf{U}^\dagger\mathbf{U} = \mathbf{1} \Leftrightarrow \mathbf{U}^\dagger = \mathbf{U}^{-1}.$$

For time-independent $\mathbf{H}$ we have

$$(\mathbf{U}(t))^{-1} = \mathbf{U}(-t),$$

that is, unitary time evolution is reversible.
Up to now we know two kinds of change of state:
i) unitary time evolution: deterministic and reversible
ii) measurement: probabilistic and irreversible.
Why is a measurement different from another physical process governed by a Hamiltonian?

1.5.4 The two-dimensional Hilbert space: Qubits and Spins

In many situations only two states of a system are important, for example, the ground and first excited states; a single spin 1/2 possesses only two states. In order to keep the analogy to classical bits as close as possible those systems are most suitable for the discussion of quantum computing.
The Hilbert space of a single spin-1/2 particle is spanned by the states

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

and

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

(The identification with $|0\rangle$ and $|1\rangle$ follows the convention of [3].) All operators in this Hilbert space 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 $\vec{B}$.
It is evident why $\mathbf{X}$ is also often called the "NOT gate" in the language of quantum computing. Any unitary 2×2 matrix is a valid quantum gate, for example the $\mathbf{Z}$ gate which generates a π relative phase. We will also frequently encounter the Hadamard gate

$$\mathbf{H} = \frac{1}{\sqrt{2}}(\mathbf{X} + \mathbf{Z}).$$

$\mathbf{H}$ (hopefully not to be confused with the Hamiltonian) is at the same time unitary and Hermitian, implying that $\mathbf{H}^2 = \mathbf{1}$. Nevertheless $\mathbf{H}$ is often called the "square root of NOT" gate, because it turns $|0\rangle$ into a state "halfway between" $|0\rangle$ and $|1\rangle$ and similarly for $|1\rangle$. (Exercise: Find the genuine square root of NOT. Hint: Try to write NOT as $\exp i\alpha\mathbf{S_x}$.)

Let us return to the spin in an external field and perform some small exercises. Consider the initial state $|\uparrow\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and determine the time evolution operator $\mathbf{U}(t)$ for a $\vec{B}$ field along one of the coordinate axes $\alpha = x, y, z$.

$$\mathbf{U}(t) = \exp\left(-\frac{i\mathbf{H}t}{\hbar}\right) = \exp\left(\frac{iB_\alpha t}{2\hbar} \underbrace{2\mathbf{S}_\alpha}_{\text{square}=\mathbf{1}}\right)$$

$$= \cos\left(\frac{B_\alpha t}{2\hbar}\right)\mathbf{1} + i\sin\left(\frac{B_\alpha t}{2\hbar}\right)2\mathbf{S}_\alpha.$$

For $\alpha = z$ we have

$$\mathbf{U}(t) = \begin{pmatrix} \exp\left(i\frac{B_\alpha t}{2\hbar}\right) & 0 \\ 0 & \exp\left(-i\frac{B_\alpha t}{2\hbar}\right) \end{pmatrix}$$

$$\Rightarrow |\psi(t)\rangle = \exp\left(i\frac{B_\alpha t}{2\hbar}\right)|\psi(0)\rangle$$

which is a stationary state, as expected, because the initial state was an eigenstate of $\mathbf{S_z}$(and thus of $\mathbf{H}$).
The case $\alpha = x$ is different:

$$\mathbf{U}(t) = \begin{pmatrix} \cos\left(\frac{B_\alpha t}{2\hbar}\right) & i\sin\left(\frac{B_\alpha t}{2\hbar}\right) \\ i\sin\left(\frac{B_\alpha t}{2\hbar}\right) & \cos\left(\frac{B_\alpha t}{2\hbar}\right) \end{pmatrix},$$

consequently

$$|\psi(t)\rangle = \begin{pmatrix} \cos\left(\frac{B_\alpha t}{2\hbar}\right) \\ i\sin\left(\frac{B_\alpha t}{2\hbar}\right) \end{pmatrix} = \cos\left(\frac{B_\alpha t}{2\hbar}\right)|\uparrow\rangle + i\sin\left(\frac{B_\alpha t}{2\hbar}\right)|\downarrow\rangle.$$

Runs through a continuum of states periodically: "uniform rotation in Hilbert space" (similar for $\alpha = y$).
A (pure) **qubit** state is an arbitrary normalized linear combination of $|\uparrow\rangle$ and $|\downarrow\rangle$ which may be parametrized, for example, by two angles:

$$|\theta, \varphi\rangle = \exp\left(-i\frac{\varphi}{2}\right)\cos\frac{\theta}{2}|\uparrow\rangle + \exp\left(i\frac{\varphi}{2}\right)\sin\frac{\theta}{2}|\downarrow\rangle$$

$$(0 \le \theta \le \pi; 0 \le \varphi \le 2\pi).$$

Thus a qubit is able to store two (bounded) *real* numbers. The question how to read, write, and manipulate this information will keep us busy for the rest of this semester (and longer, maybe).
It is easy to check that the above qubit is an eigenstate of the operator

$$\cos\theta\mathbf{S_z} + \sin\theta\cos\varphi\mathbf{S_x} + \sin\theta\sin\varphi\mathbf{S_y}$$

with eigenvalue $+1/2$. Thus, in order to prepare this qubit, one "only" needs to align the spin along the (θ, φ) direction by a sufficiently strong field.

1.5.5 Two qubits and the density matrix

"Much that is weird and wonderful about quantum mechanics can be appreciated by considering the properties of the quantum states of two qubits." (John Preskill [7])
In the real world there are no isolated spins 1/2; quantum systems always couple to the "environment" which we often cannot or do not want to take into account in our quantum mechanical considerations. However, if we consider a quantum system which is in reality only part of a larger system, we will have to abandon some of our dogmas, because

- states are no longer vectors in Hilbert space,
- measurements are no longer orthogonal projections on to the final state
- and time evolution is no longer unitary.

The simplest example is given by one qubit A = "system" (accessible) and another qubit B = "environment" (inaccessible). $\{|\uparrow\rangle_A, |\downarrow\rangle_A\}$ and $\{|\uparrow\rangle_B, |\downarrow\rangle_B\}$ are orthonormal bases for the two subsystems.
A **correlated** or **entangled** state of two qubits:

$$|\psi\rangle = a\,|\uparrow\rangle_A \otimes |\uparrow\rangle_B + b\,|\downarrow\rangle_A \otimes |\downarrow\rangle_B.$$

A measurement of the state of qubit A (that is, a projection on to the A basis) yields $|\uparrow\rangle_A \otimes |\uparrow\rangle_B$ with probability $|a|^2$ and $|\downarrow\rangle_A \otimes |\downarrow\rangle_B$ with probability $|b|^2$. In both cases after the measurement on A the state of B is *fixed*.
Now measure an observable which acts only on A:

$$\mathbf{M}_A \otimes \mathbf{1}_B.$$

Expectation value of this observable in the state $|\psi\rangle$:

$$\langle \mathbf{M}_A \rangle = \langle \psi | \mathbf{M}_A \otimes \mathbf{1}_B | \psi \rangle =$$

$$[a^* \, {}_A\langle\uparrow| \otimes_B \langle\uparrow| + b^* \, {}_A\langle\downarrow| \otimes_B \langle\downarrow|] \, \mathbf{M}_A \otimes \mathbf{1}_B \, [a \, |\uparrow\rangle_A \otimes |\uparrow\rangle_B + b \, |\downarrow\rangle_A \otimes |\downarrow\rangle_B] = \cdots$$

$\mathbf{1}_B$ does *not* change $|...\rangle_B$ states; ${}_B\langle\uparrow|\downarrow\rangle_B = 0 \Rightarrow$ only two terms survive.

$$\cdots = |a|^2 \, {}_A\langle\uparrow| \mathbf{M}_A |\uparrow\rangle_A + |b|^2 \, {}_A\langle\downarrow| \mathbf{M}_A |\downarrow\rangle_A = \mathrm{Tr}\left(|a|^2 \, \mathbf{P}_{\uparrow A} \, \mathbf{M}_A + |b|^2 \, \mathbf{P}_{\downarrow A} \, \mathbf{M}_A \right) =$$

$$= \mathrm{Tr}\left(\left[|a|^2 \, \mathbf{P}_{\uparrow A} + |b|^2 \, \mathbf{P}_{\downarrow A} \right] \mathbf{M}_A \right) = \mathrm{Tr}\left(\rho_A \mathbf{M}_A \right).$$

Here

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

are projectors on to the $\uparrow, \downarrow$ subspaces of A; Tr denotes the trace (sum of the diagonal elements) in the Hilbert space of A, that is,

$$\mathrm{Tr}\, \mathbf{X} = {}_A\langle\uparrow| \mathbf{X} |\uparrow\rangle_A + {}_A\langle\downarrow| \mathbf{X} |\downarrow\rangle_A.$$

The quantity

$$\rho_A = |a|^2 \mathbf{P}_{\uparrow A} + |b|^2 \mathbf{P}_{\downarrow A} = \begin{pmatrix} |a|^2 & 0 \\ 0 & |b|^2 \end{pmatrix}$$

is the **density operator** (density matrix); it is Hermitian, positive and its trace is unity. It is important to note that *every* operator with these properties is a possible density operator, irrespective of its being diagonal or not in the basis which we have chosen accidentally or thoughtfully!
If $\rho_A^2 = \rho_A$ (for example if $|a| = 1$ in our example) ρ_A is a projector on a vector in Hilbert space: *pure* state; otherwise: *mixed* state; often called "incoherent superposition" by people who like optics better than quantum mechanics. The above ρ_A is obviously a mixed state.
The initial $|\psi\rangle$ above was a vector in the "large" Hilbert space; the two subsystems were correlated (entangled). Let us now consider an *uncorrelated* state:

$$|\Phi\rangle = (a|\uparrow\rangle_A + b|\downarrow\rangle_A) \otimes (c|\uparrow\rangle_B + d|\downarrow\rangle_B)$$

with $|a|^2 + |b|^2 = |c|^2 + |d|^2 = 1$. For that state we quickly end up with

$$\langle \mathbf{M_A} \rangle = \langle \Phi | \mathbf{M_A} \otimes \mathbf{1_B} | \Phi \rangle$$

$$= [a^* \, {}_A\langle\uparrow| + b^* \, {}_A\langle\downarrow|] \, \mathbf{M_A} \, [a \, |\uparrow\rangle_A + b \, |\downarrow\rangle_A] \, [c^* \, {}_B\langle\uparrow| + d^* \, {}_B\langle\downarrow|] \, \mathbf{1_B} \, [c \, |\uparrow\rangle_B + d \, |\downarrow\rangle_B].$$

The "B part" of course yields unity and we obtain

$$\langle \mathbf{M_A} \rangle = |a|^2 \, {}_A\langle\uparrow| \mathbf{M_A} |\uparrow\rangle_A + |b|^2 \, {}_A\langle\downarrow| \mathbf{M_A} |\downarrow\rangle_A + a^* b \, {}_A\langle\uparrow| \mathbf{M_A} |\downarrow\rangle_A + b^* a \, {}_A\langle\downarrow| \mathbf{M_A} |\uparrow\rangle$$

$$= \mathrm{Tr}\left(\mathbf{M_A} \left[|a|^2 \mathbf{P}_{\uparrow \mathbf{A}} + |b|^2 \mathbf{P}_{\downarrow \mathbf{A}} + a^* b |\downarrow\rangle_{AA}\langle\uparrow| + b^* a |\uparrow\rangle_{AA}\langle\downarrow| \right] \right) = \mathrm{Tr}\left(\mathbf{M_A} \rho_{\mathbf{A}} \right).$$

Again $\rho_{\mathbf{A}}$ is Hermitian and of unit trace, but obviously *not* diagonal; in the usual basis it is

$$\rho_{\mathbf{A}} = \begin{pmatrix} |a|^2 & a^* b \\ b^* a & |b|^2 \end{pmatrix}.$$

Nevertheless ${\rho_{\mathbf{A}}}^2 = \rho_{\mathbf{A}}$, as we can easily calculate: the density matrix of A is a pure state if the initial (pure) state of the combined system $A+B$ is a product state (that is, uncorrelated, not entangled). If the (pure) state of the combined system $A + B$ is entangled, the summation over all possibilities for the state of B ("partial trace over the Hilbert space of B") leads to the loss of the phases of a and b and we end up with a mixed state. The following picture for the loss of coherence thus arises: in the beginning, system and environment are uncorrelated. The system's density matrix is initially pure. By interaction, system and environment become entangled (we will soon see examples for this) and the system's density matrix becomes mixed.

The general density matrix of a qubit:

$$\rho(\vec{P}) = \frac{1}{2}\left(\mathbf{1} + 2\vec{P}\cdot\tilde{\mathbf{S}}\right); \quad |\vec{P}| \leq 1$$

The possible $\vec{P}$ form the Bloch sphere; pure states have $|\vec{P}| = 1$.
Exercise: Distinguish the pure state

$$|\chi\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)$$

from the mixed state

$$\rho = \frac{1}{2}(\mathbf{P}_\uparrow + \mathbf{P}_\downarrow)$$

by determining the expectation values and variances of the operators $\mathbf{S}_\alpha$. Furthermore, determine the polarization vector $\vec{P}$ in both cases.
Further important theoretical concepts which could be discussed at this point are

- the time evolution of the density matrix (generalization from the Schrödinger equation to the Liouville / von Neumann equation)
- more formal definition of entanglement.

The first item is treated in many quantum mechanics texts, whereas the second one is still a topic of research, see [7].