Quantum Information Processing — Quanteninformationsverarbeitung

Joachim Stolze and Dieter Suter
Institut für Physik, Universität Dortmund
SS 2002
(Dated: May 7, 2002)

A two-hour introductory course mainly designed for the graduate students of the "Graduiertenkolleg 726: Materialeigenschaften und Konzepte für die Quanteninformationsverarbeitung".

I. INTRODUCTION AND SURVEY

A. Literature

General references for this course are the books [1, 2, 3, 4] and also the review articles [5, 6, 7]. 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 [8] 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.

B. Motivation: why quantum computers

Many different reasons exist to investigate quantum computers. For most of us, the main motivation is probably the fascinating physics that coalesces in this field. This motivation is very strong for people that are already actively doing research in a related area, but it is rather hard to communicate to people outside physics. On the other hand, the speed at which a field evolves is often tied to the successful communication to outsiders of the prospects for the field. One aspect of this is certainly the possible economic implications of this emerging technology.

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

The feature size of electronic devices is now in the range of 100 nm and decreasing at a rate of some 12% per year. According 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 $5 \cdot 10^{-18} F$. A change in the voltage of 0.1 V will then move less than four electrons in such a device.

This makes it obvious that the progress that we have today will soon lead to a situation where it is no longer possible to describe the flow of electricity as a classical current. While a classical device, such as the workhorse FET, requires a continuous relation between current and voltage, this will no longer be the case in the quantum mechanical regime, as experimental prototypes clearly show.

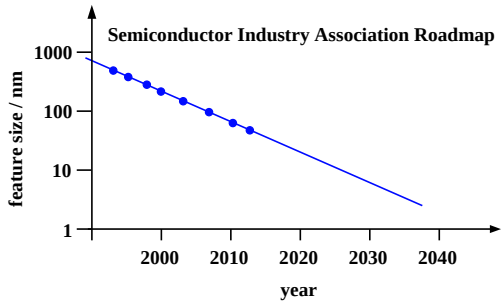

FIG. I.1: Prospective evolution of feature size in microelectronic circuits (source: international semiconductor association roadmap).

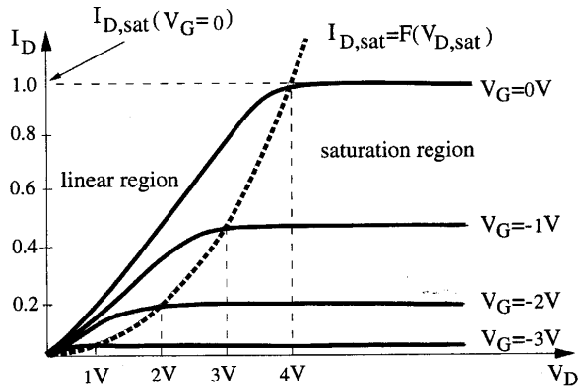

FIG. I.2: Current/voltage characteristics of classical FET.

As an example, the figure shows the behavior of an experimental device where the transfer of single electrons is observed.

Possibly even more impressive is a consideration of the energy dissipated in a logical step. Over the last fifty years, this number has decreased by more than ten orders of magnitude, again following an exponential time dependence. A straightforward extrapolation shows that this trend would decrease the dissipated energy to less than kT in little more than ten

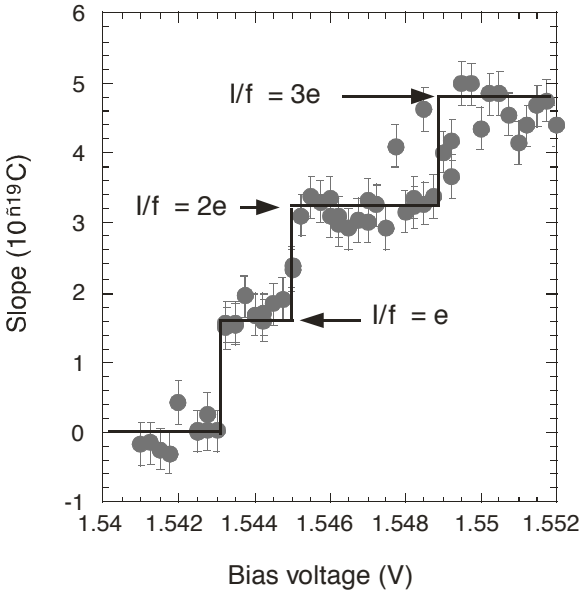

FIG. I.3: Behavior of a single electron device.

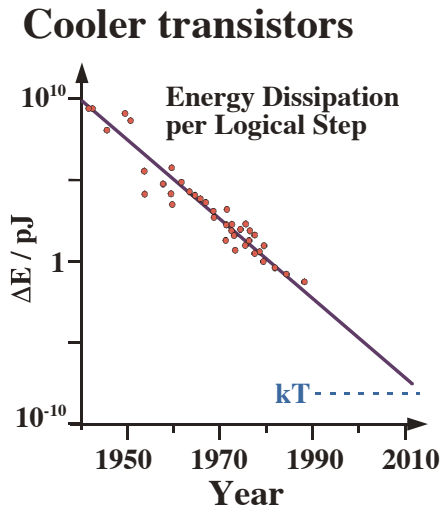

FIG. I.4: Energy dissipation of logic devices vs. time.

years. This implies that around this time, switches would start to become unstable, since thermal fluctuations would be sufficient to trigger them. Standard electronic circuitry would then no longer be able to function.

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 studies the possibilities provided by such an approach and tries to develop algorithms and implementations to harness this potential.

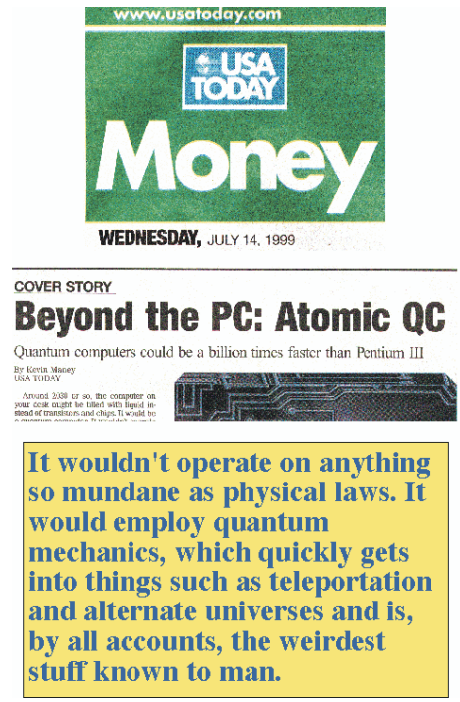

FIG. I.5: Quantum computers in the media.

The possibility that this emerging technology will lead to significantly more powerful computers has generated a lot of attention, even outside the scientific community. In particular when the first experimental results were published, this "revolutionary new type of computers" reached the headlines of many news media. Of course many journalists were not able to grasp all the essentials of this approach.

C. Some fundamental questions

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 partic-

ular, logical gates like AND are clearly not reversible, since it is not possible to calculate the input from the output. When quantum mechanics is used to describe the operation of a computer, the computational process, like all other time-dependent processes, is just an evolution under a suitable Hamiltonian. Quantum mechanics, on the other hand, describes all time evolution with unitary operators, such as $\exp(iHt)$, which are obviously time-reversible ($\exp(iHt)\exp(-iHt) = 1$).

D. History of quantum information processing

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 [9] 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 [10].

1982 Richard Feynman

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

The computational power required to simulate quantum mechanical systems grows exponentially with the size of the system

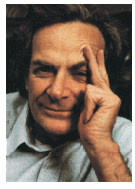

Efficient simulation requires a quantum mechanical computer.

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. Feynman 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.

E. Shor's factoring algorithm

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 [11].

1994 Peter Shor

P. Shor, *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*, in *35th Annual Symposium on Foundations of Computer Science*, IEEE Press, Piscataway, NJ (1994).

With his factoring algorithm, the computation time grows only algebraically, rather than exponentially with the number of digits.

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 implies that for a sufficiently large number, a quantum computer will always finish the factorization faster than a classical computer, even if the classical computer runs on a much faster clock.
We illustrate this with a numerical 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. 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.

F. Quantum information

Classically, information is encoded in a sequence of bits, i.e. entities which can be in two distinguishable states, which are conventionally labeled with 0 and 1.

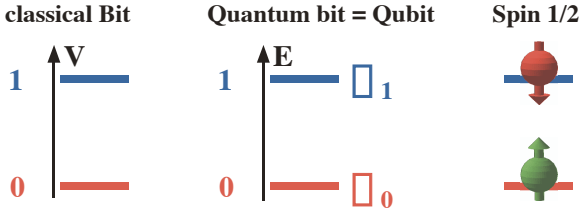

Quantum mechanically, these entities are quantum mechanical systems with (at least) two distinguishable states ψ_0 andψ_1. A typical example is a spin 1/2, which has two possible states. Another example is a photon, which can be polarized either vertically or horizontally. One of these states is identified with the logical value 0 (or false), the other with the value 1 (or true).
The main difference between quantum mechanical and classical information is that in the quantum mechanical case, the system is not necessarily in the state 0 or 1. Instead it can be in an arbitrary superposition (linear combination) of these states. The power of quantum computers is directly related to this possibility of creating superpositions of states and applying logical operations to them: this allows one to perform many operations in parallel. For N qubits, it is possible to create a superposition of 2^N states. Logical operations like multiplications can then be applied to all those states simultaneously.
Becoming slightly more formal, we find that the information, which is encoded in a quantum mechanical system (or quantum register), is described by a vector in Hilbert space. For the simplest case of a single qubit, the state is $\psi = a\psi_0 + b\psi_1$. The two parameters a and b are both complex numbers. Taking normalization into account, the system is therefore described by three continuous variables.
The fact that the state is described by three continuous variables does not imply that a single photon can store an infinite amount of information. To obtain the information content, one has to take the measurement process into account, which retrieves the information: 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. Without a complete analysis, this can be rationalized by the consideration that we can make two independent measurements on a photon, corresponding, e.g., to the measurement of the polarization horizontal/vertical or at ± 45 degrees.
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.

G. No cloning theorem

Since the no-cloning theorem is so fundamental to quantum information processing, we briefly repeat here its proof, which is due to Wootters and Zurek [12] and also Dieks [13].
Let $|\psi\rangle$ be a pure state, and $|s\rangle$ some "standard" (initial) state. A "quantum state cloner" would then turn the standard state into an exact copy of the original, without affecting $|\psi\rangle$. Mathematically, this could be described by a unitary operator $\mathbf{U}$ with

$$\mathbf{U}|\psi\rangle|s\rangle = |\psi\rangle|\psi\rangle \quad \forall|\psi\rangle.$$

Now consider the cloning of a second state $|\phi\rangle$:

$$\mathbf{U}|\phi\rangle|s\rangle = |\phi\rangle|\phi\rangle.$$

For simplicity assume that $|\psi\rangle$, $|\phi\rangle$, and $|s\rangle$ are normalized. Take the scalar product of the two equations above and keep in mind that $\mathbf{U}$ is unitary, that is, it preserves scalar products:

$$\left(\langle s|\langle\psi|\mathbf{U}^\dagger\right)\left(\mathbf{U}|\phi\rangle|s\rangle\right) = \langle s|s\rangle\langle\psi|\phi\rangle = \langle\psi|\phi\rangle$$

$$= \left(\langle\psi|\langle\psi|\right)\left(|\phi\rangle|\phi\rangle\right) = \left(\langle\psi|\phi\rangle\right)^2$$

and this is possible only if $\langle\psi|\phi\rangle = 0$ or $\langle\psi|\phi\rangle = 1$, that is, if the two states to be copied by the same operation are either identical or orthogonal. This is in contradiction to our requirement that the cloning should work for all states ψ and proves that general cloning is not possible.

H. Quantum 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. This is typically done by encoding the information in photons. 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 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.
This is not automatic, however. 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. changing the number of photons). Heisenberg's principle requires, however, that such a measurement affects the conjugate variable, in this example the phase of the photon.

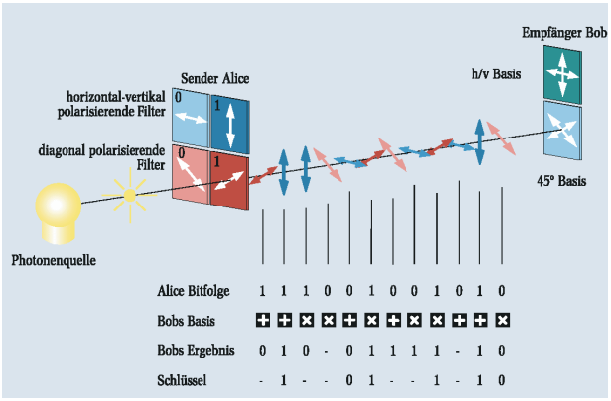

FIG. I.6: Possible protocol for quantum key distribution.

The two partners can use this fact to make the communication protocol secure. A typical protocol requires one of the two partners (typically called Alice) to send a stream of photons to the second partner (typically called Bob), which are entangled with a second set of photons, which Alice keeps. The two partners then make a measurement of the polarization of these photons, choosing the axis of the polarizer with a random number generator. If the photon pairs are originally in a singlet state, each partner knows then the result of the other partner's measurements provided they used the same axis of the polarizer. They can therefore generate a common secret string of bits by exchanging through a public channel (e.g. a radio transmission) the orientation of the polarizer that they used for their measurements. This scheme has been tested in a number of experiments by using optical fibers or beams through free space.

End of week 1, April 19.

I. 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 and can therefore be described by the spinor

$$|\psi\rangle = c_0|0\rangle + c_1|1\rangle.$$

The total collection of qubits is called a quantum register. Its state is written as

$$|\psi\rangle^{reg} = c_0|0,0,0..0\rangle + c_1|0,0,0..1\rangle + c_2|0,0,0..1,0\rangle + ...$$

While today's quantum registers are limited to 7 qubits, a useful quantum computer will require several hundred to 1000 qubits.
Before an actual computation can be initiated, the quantum register must be initialized into a well defined state, typically the quantum mechanical ground state $|0,0,...0>$. This operation is non-unitary, i.e. it always must include dissipation.
The actual information processing occurs through the operation of quantum gates that operate on the quantum register. The sequence of quantum gates required depends on the specific algorithm. The program that specifies this sequence may be stored in a classical device associated with the quantum computer, such as a classical computer.
The logic operations that can be performed on a quantum computer correspond to changes in the information stored in the quantum register. While the quantum register is represented by a state vector, the operations applied to it are 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 $R(\alpha;\theta,\phi)$ by an angle α around a rotation axis whose orientation is specified by the polar angles θ,ϕ.

- 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	01
1	0	11
1	1	10

If the control bit is 1, the target bit is switched, otherwise it is not changed. The control bit is never changed.
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 thus 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 2-qubit operations is larger, but increases only linearly with the number of qubits. The overall process therefore remains efficient.

J. Decoherence

Possibly the biggest obstacle to overcome when one tries to build a quantum computer is decoherence. This term summarizes all processes that contribute to a decay of the information coded in the quantum register. As we have stressed above, the power of quantum computers derives from the possibility to perform logical operations on a large number of states simultaneously, which have been combined into a superposition state. If the relative phase between these states slips, the result of the operation will effectively become associated with the wrong input, thereby destroying the information. The larger the number of qubits, the more powerful the processor becomes, but the more fragile the quantum information gets.
The biggest contribution to decoherence is usually dephasing. To illustrate its effect, we consider a system that is initially in a superposition state

$$|\psi(0)\rangle = a|0\rangle + b|1\rangle.$$

If the two states $|0\rangle$ and $|1\rangle$ are eigenstates of the driving Hamiltonian H with eigenvalues E_0 and E_1, an ideal evolution will transform this state into

$$|\psi(t)\rangle = a|0\rangle e^{iE_0 t/\hbar} + b|1\rangle e^{iE_1 t/\hbar}.$$

Dephasing is due to additional (uncontrollable) interactions, which change the relative phase between the states by an angle δ

$$|\psi(t)\rangle = a|0\rangle e^{iE_0 t/\hbar} e^{i\delta/2} + b|1\rangle e^{iE_1 t/\hbar} e^{-i\delta/2}.$$

In a simple picture, this additional phase can be represented as an additional rotation of the state vector. On a purely classical basis, this additional phase increment may arise from experimental uncertainties, such as logical gate errors. This will then lead to different evolutions of specific implementations. Every specific quantum computer remains then in a coherent superposition, although this may not be exactly the state that would correspond to the mathematically correct representation.

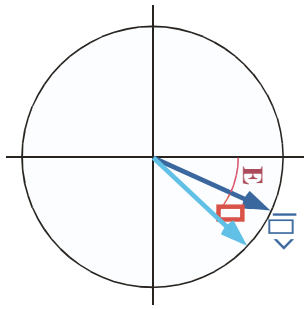

FIG. I.7: Coherent and incoherent contribution to the evolution.

In quantum mechanics, the situation is slightly different. Here, these phase-kicks are correlated to states of the external system, which is referred to as the bath. Typical examples for relevant degrees of freedom in the environment are phonons passing through the system or photons of the background radiation. For every state of this external system, the quantum register remains in a pure state, but the phase δ for this realization will be different than for other states of the environment.

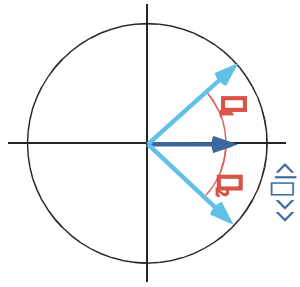

FIG. I.8: Mixed state from ensemble average.

Since it is never possible to exactly know the state of the external system, one has to average over all accessible states of the external system. This averaging process changes the situation qualitatively: the vector representing the system is no longer only rotated by these additional phase kicks, it also becomes shorter. Technically, it is no longer in a pure state, but rather in a mixed state. In the simple picture given above, the vector no longer ends on the unit circle (or sphere), but remains inside it.
In most systems, an exact description of these processes is not available. One then usually approximates the effect by an exponential decay process for the relevant density operator elements. For the off-diagonal elements one writes

$$\rho_{ij}(t) = \rho_{ij}(0) e^{i(E_i - E_j)t/\hbar} e^{-t/T_2},$$

while the diagonal elements approach thermal equilibrium with a time constant T_1. These longitudinal relaxation processes also affect the quantum computation, causing a decay of the information. However, they are also needed, since they bring the system to the ground state, as required for initialization.

The effect of decoherence is a loss of information in the system. Since it is highly unlikely that any system will be able to successfully complete a useful quantum information process before decoherence occurs, it is vital to develop strategies that eliminate or reduce the effect of decoherence. One possibility that is pursued actively, is to apply quantum error correction. Basically these processes use coding of quantum information in additional qubits and intermittently check which of them have changed; suitable algorithms have been decveloped for retrieving the lost information from the excess qubits.

K. Implementations

To actually build a quantum computer, a suitable physical system has to be identified and the associated controls must be put in place. We give here a brief overview of the conditions that implementations must fulfill and discuss some issues that help in identifying suitable systems.

The quantum information is stored in a register. Any implementation therefore has to define a quantum mechanical system that provides the quantum register containing N qubits. For a "useful" quantum computer, N should be at least 400, better 1000; limitations on the number N of identifiable qubits will therefore be an important consideration.

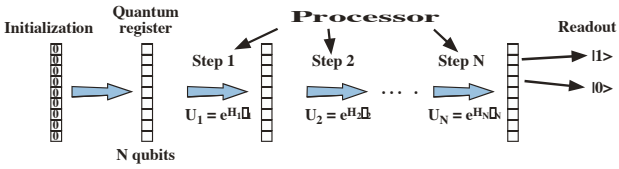

FIG. I.9: Principle of operation of quantum processor.

These qubits must be initialized into a well defined state, typically into a ground state $|0\rangle$. This is necessarily a dissipative process. Implementations must therefore provide a suitable mechanism for initialization.

The implementation must then provide a mechanism for applying computational steps to the quantum register. Each of these steps consists of a unitary operation $e^{-iH_i\tau_i}$ defined by a Hamiltonian H_i that is applied for a time τ_i. The Hamiltonian must act on specific qubits and pairs of qubits by applying electromagnetic fields. The quantum computer must therefore contain mechanisms for generating these fields in a well controlled manner. The details of the requirements on these fields depend on the system chosen for the quantum register.

After the last processing step, the resulting state of the quantum register must be determined, i.e. the result of the computation must be read out. This would typically correspond to an ideal quantum mechanical measurement, i.e. the projection onto the eigenstate of the corresponding observable. Readout has to be done on each qubit separately.

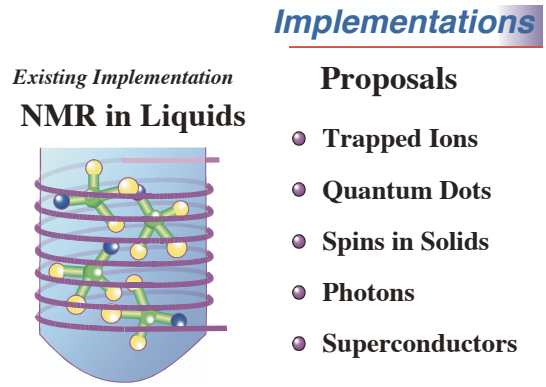

FIG. I.10: Existing and proposed implementations of quantum omputers.

Today, a single implementation of a quantum computer exists, which stores the qubits in nuclear spin states of molecules in solution, i.e. liquid-state NMR. Details of this implementation have been discussed during the last semester.

In addition to this single existing implementation, there is a long list of proposed implementations, which includes, as qubits, nuclear and electronic spins, photons, trapped ions, as well as various states of quantum confined structures, mostly in semiconductors, and superconducting devices such as Josephson junctions. Most projects within this "Graduiertenkolleg" try to assess and improve the suitability of some of these proposals. For some of these systems, operations have been demonstrated that may be considered one-qubit operations. However, none of them has demonstrated two-qubit operations. The implementation of a full quantum algorithm is therefore in most cases still in the distant future. The one system that is most likely to implement a simple quantum algorithm is the trapped ion system.

L. DiVincenzo's criteria

DiVincenzo [14] gives five requirements that a quantum computer must fulfill:

1) A scalable physical system with well characterized qubits. An implementation or embodiment of qubits corresponds to a physical system that has at least two energy levels that can be identified with the two logical states $|0\rangle$ and $|1\rangle$. The system will always be characterized by a number of parameters. They determine the internal state of the qubit, i.e. the energies of the states $|0\rangle$ and $|1\rangle$, but also of any additional states that may exist in the system being considered. All of these parameters must be well known to provide for the precise implementation of quantum gates.

The system chosen for storing the qubits must allow the creation of arbitrary superpositions of the basis states. This is usually possible unless there is a selec-

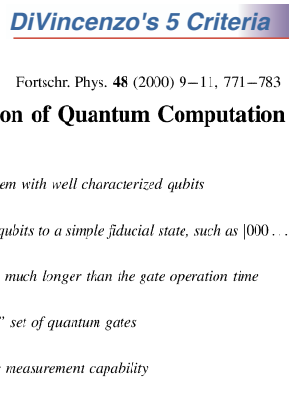

FIG. I.11: The five criteria, put forward by di Vincenzo, that a quantum computer must fulfill.

tion rule that prevents it. As an example, we consider two neighbouring quantum dots, where an electron can tunnel from one dot to the other. It is then possible to identify state $|0\rangle$ with the electron being in dot 1 and state $|1\rangle$ with the electron being in dot 2. However, it is not possible to identify a qubit with a single quantum dot, e.g. with the assignment that the presence of an electron corresponds to $|1\rangle$, while its absence would correspond to $|0\rangle$. The superposition of these two states would then correspond to a superposition between states with different particle numbers, which is usually impossible to achieve.

Besides the internal Hamiltonian, the interaction of the system with external fields is also important. External fields are generally required to apply logical operations to the qubits. Finally, the couplings between different qubits must be described, as they are needed for logical operations.

DiVincenzo's second requirement is

2) Initialization into a well defined state. Typically, this state is chosen equal to the logical state $|0\rangle$ for all qubits. In principle thermal relaxation may achieve this, provided that the thermal energy k_BT is small compared to the energy level splitting between states $|0\rangle$ and $|1\rangle$. This may be a slow process in many systems, in particular in the spin systems, where the relaxation times are long. This is not critical for the computation process itself; however, future quantum computers will require error correction schemes. All error correction schemes known to date require an input in the form of freshly initialized qubits. These error correction qubits must be intialized at a rate that is fast compared to the dephasing rate. This requirement cannot be fulfilled by thermal relaxation, where the dephasing processes are always faster than the spin-lattice relaxation. The requirement can be met, however, in many optical systems, such as ion traps, where the initialization goes through optical excitation, which may proceed over a time of the order of nanoseconds.

3) Long decoherence times. The information in the quantum register is subject to decay, due to the interaction with external degrees of freedom. The computation must therefore be completed before this decay has significantly degraded the information. The relevant figure of merit is the number of gate operations that can be completed before a decoherence time.

The effect of decoherence can partly be eliminated by quantum error correction. However, error correction also increases the duration of the computation and introduces additional errors. It has been shown that computation can proceed for an arbitrary duration if quantum error correction is used and error-free computation without error correction is possible for a critical minimum duration that is of the order of some tens of thousands of gate operations.

4) A universal set of quantum gates. The unitary operations that act as gates on the qubits must be implemented by Hamiltonians that act on the system for a specified time. Generating the single qubits Hamiltonians is in general relatively straightforward: typically they correspond to external fields acting on the qubits for a spcified duration. More complicated are the 2-qubit operations, which cannot be implemented by external fields alone. They involve interactions between the qubits, and in many cases these interactions cannot be switched on and off. Often one has to use static interactions and eliminate the unwanted ones by a procedure called refocusing.

Every experimentally realisable gate will include imperfections, i.e. deviations from the ideal behavior. This has the effect of degrading the information in the quantum register and is therefore similar to an additional source of decoherence. Consequently, these errors can also be eliminated by error correction schemes, provided they are small enough.

5) A qubit-selective readout. Such a readout represents a measurement in the quantum mechanical sense. An ideal quantum mechanical measurement collapses the state ψ into an eigenstate ϕ_i of the observable and returns the eigenvalue λ_i of the corresponding state with probability $|c_i|^2$, where c_i is the expansion coefficient of the state $\psi = \sum c_i\phi_i$.

Real measurements deviate from this. In many realistic systems, measurement attempts will return no result instead, e.g. when one tries to measure the state of a qubit by scattering a photon from it. If the photon is not scattered, this is not important, one just repeats the attempt. If the photon is scattered but not detected, this is more critical: In this case, an interaction of the qubit with an external system (the photon) has changed the state of the qubit, and a repetition of the measurements may produce a different result.

Several strategies are possible to circumvent this problem: one can try to use a QND (=quantum nondemolition measurement). Such a measurement arranges for the unavoidable influence that the measurement must have on the qubit to be such that it does not affect later measurements of the same variable. Not all variables can be measured this way, but in most cases it should be possible to arrange the system in such a way that QND measurements can be used at least in principle.

Another possibility is to read out not the qubit itself, but a copy of it. If the measurement is not successful, or to check the validity of the measurement result, one can then make an additional copy and read that out. Such a procedure could be repeated many times to achieve very reliable readout even with very unreliable single measurements. The critical part here is the copy operation, which must be reliable. As we have stressed before, it is not possible to clone a quantum mechanical state, i.e. to make a perfect copy. However, copying just the information of a quantum mechanical state that is relevant for the readout of a specific variable is perfectly possible (in principle!) and can be repeated arbitrarily often.