

1 Introduction and survey

1.1 Information, computers and quantum mechanics

1.1.1 History of computing

Computing is based on numbers. The concept of numbers and counting emerged probably some 40000 years ago. Actual systems for representing arbitrary numbers were developed around 5000 years ago in Mesopotamia and led to the first mathematical procedures like arithmetics, algebra and geometry. These concepts were motivated by applications in taxation, commerce, trade, astronomy and calendars. In Egypt, e.g., predicting the flooding of the Nile was essential for agriculture.

More abstract mathematical concepts were developed in Greece by philosophers like Pythagoras and Euclid. The first formal collection of algorithms (computational recipes) was written by Abu Dscha'far Muhammad ibn Musa al-Chwarizmi (780-850).

Some tasks were aided by computational devices, which could be analog (like the slide rule) or digital (like the abacus).

1.1.2 General purpose digital computers

Analog computers have been used for thousands of years, mostly to calculate specific events like the motion of planets. Digital computers like the abacus were used mostly for simple algebraic operations like adding numbers. More sophisticated models were designed and built in the 17th century by Wilhelm Schickard, Blaise Pascal and others.¹

¹U useful summary of important contributions to the history of computation can be found at




Figure 1.1: Two types of digital computers: the abacus (left) and the Zuse Z1 (right).

In the 19th century, Charles Babbage combined such calculators with programmable input from a reading device. His machine never achieved its targeted function, but around 1940, the German engineer Konrad Zuse built the first freely programmable prototype computers that can be proved to be general-purpose: they can be used to solve any computable problem in finite time. The right-hand picture in figure 1.1 shows the first of these devices.

During the following years, the field moved from mechanical to electronic devices, using first vacuum tubes and then solid-state electronics. These devices soon exceeded the computing capabilities of humans and mechanical computers, while also getting cheaper thanks to economy of scale. The combination of these factors with others, in particular communication resulted in a transformation of the economy and culture that is often called the "Information Age" or the 3rd industrial revolution.

1.1.3 The digital revolution

Storage, interchange and processing of information is a defining feature of human culture as well

https://en.wikipedia.org/wiki/List_of_pioneers_in_computer_science
https://en.wikipedia.org/wiki/List_of_pioneers_in_computer_science

as the basis of our economic system. Over the last fifty years, all these processes have undergone dramatic changes, driven by the evolution of microelectronics technology. The increasing availability of cheap storage, fast processors and global telecommunication (including the internet) has prompted a shift from a number of different conventional techniques used to store, process and transmit information, which used different, mostly analog techniques, to those which use all-digital forms of representing information.

This convergence of technologies has also eased the connection between storage, processing and communication and made most of the ongoing processes transparent or invisible to the person who is actually using them. A search for a picture over an internet search engine, e.g., which typically involves typing a few words and results in a long list of "hits", involves all three types of processes mentioned several times:

- The computer on which the person works interprets the input and uses its locally stored information to decide what action it has to take.
- It communicates with routers to obtain the address of the search engine.
- It sends the request over the internet to the search engine. The transfer of information over the internet involves multiple steps of processing and using stored information about connections at all nodes.
- The search engine receives the request and compares the keywords to those stored in its files.
- It uses stored rules to rank the hits.
- The result is sent back over the internet.
- The workstation receives the information and uses stored information to display the information.

Each of these steps can be further subdivided into smaller steps that may again include different types of actions on the information being exchanged between many different parties (most of them electronic circuits).

These fundamental changes of the way in which information is represented and processed have simultaneously changed the way in which we use information. One consequence is that, very often, information can no longer be localized or associated with a specific physical device. While hand-written notes and symbols engraved in rocks represented unique instances of the pertinent information, every electronic message is stored (at least temporarily) on many different computers. It is therefore not only available to the person to whom it is sent and to the person who wrote it, but also to many others like system managers, hackers, or government agencies.

Most users of digital information experience the paradigm shift from conventional forms of information representation to a unified digital form as an exciting possibility for improved communication, easier access to vital information and additional choices for entertainment. This attitude has driven the growth of the microelectronics industry over the last decades and is likely to remain an important economic force for the foreseeable future.

At the same time, the global availability of information and the difficulty of controlling one's personal data have prompted concerns about maintaining privacy. The emerging field of quantum information processing holds promises that are relevant for both issues, the further evolution of microelectronics as well as the concerns about privacy. This field, which combines approaches from physics, mathematics, and computer science, differs from conventional approaches by taking into account the quantum mechanical nature of the physical devices that store and process the information. In this course, we concentrate on the aspect of "quantum computers", which refers to machines built on the basis of explicitly quantum mechanical systems and designed to process information in a way that is much more efficient than conventional computers. While it is still unclear at what time (and if ever) such computers will be more powerful than classical computers, it is quite clear that at least some of the underlying physics will be in-

corporated into future generations of information processing hardware. The related field of quantum communication, which promises to deliver ways of exchanging information that cannot be tapped by any eavesdropper, will only be mentioned briefly in section 13.

1.1.4 Moore's law

The evolution of micro- and optoelectronic devices and the associated digitization of information has relied on improvements in the fabrication of semiconductors that have led to ever smaller and faster components. The decrease in size, in particular, has allowed more components to be packed onto a chip, thus making them more powerful by integrating more functions. Simultaneously, the decrease in size is a prerequisite for making faster devices, as long as they rely on a fixed, systemwide clock. As early as 1965, Gordon Moore noticed that the number of components that could be placed on a chip had grown exponentially over many years, while the feature size had shrunk at a similar rate 1. This trend continued over the next forty years and is expected to do so for several more years.




Figure 1.2: Prospective evolution of feature size in microelectronic circuits (source: international semiconductor association roadmap).

Figure 1.2 shows the current expectations: it

represents the projections that the semiconductor industry association makes for the coming decade. As shown in Fig. 1.2, the feature size of electronic devices is now less than 10 nm and decreasing at a rate of some 12% per year.

This trend could in principle continue for another twenty years before the ultimate limit is reached, which corresponds to the size of an atom. TSMC began production of its 2 nm process in July 2024, with mass production planned for the second half of 2025, and Samsung plans to start production in 2025. However, the "2 nm" node does not relate to actual physical sizes like gate length, metal pitch or gate pitch. The actual feature sizes are of the order of 20-50 nm.

Much before the ultimate limit of single atoms, the feature size will become smaller than some less well defined limit, where the electrons that do the work in the semiconductor devices, will start to show that their behavior is governed by quantum mechanics, rather than the classical physical laws that are currently used to describe their behavior.

1.1.5 Emergence of quantum behavior

The reduction of feature size also implies a decrease in operation voltage, since the internal fields would otherwise exceed the breakthrough fields of all available materials. Within the next ten years, the operational voltage is expected to decrease to less than one Volt.

Figure 1.3 shows how the transition to the quantum regime will change the way in which typical electronic devices operate. Capacitors, which are present in many electronic circuits, exhibit a direct proportionality between applied voltage and stored charge in all classical devices. The capacitance of a spherical capacitor is $C = 4\pi\epsilon_0 r$. For a spherical capacitor with radius 5 nm, the capacitance is therefore of the order of $5.5 \cdot 10^{-19}$ F. A change in the voltage of 1 V will then move less than four electrons in such a device, again making quantization effects noticeable. When the capacitance becomes small enough, the transfer




Figure 1.3: Current/voltage characteristics of classical capacitor (left) and its analog in the quantum regime, where individual electrons can or cannot enter the device.

of a single electron will change the potential of the capacitor by a large enough amount that it takes a significantly larger voltage to transfer additional charges. While the capacitance of real capacitors is higher, the number of electrons stored in a memory cell will become a small integer number in the near future, again bringing quantum physics into play.

Classical physics is an approximation of the more fundamental laws of quantum mechanics, which, until today, has proved sufficiently accurate for all fields of engineering. Quantum mechanics is required in order to understand the properties of all types of matter, such as the distinction between insulators, semiconductors and metals. It allows one to understand the properties of semiconductors, such as current-voltage curves of diodes, from their microscopic structure. Once these properties are established, however, it becomes possible to describe the operation of semiconductor devices on the basis of the classical theory of electrodynamics.

This classical description of the operation of semiconductor devices will become impossible when the feature size reaches the coherence length. This quantity depends on the details of the material, the processing and the temperature at which the device operates, but typically is in the range of a few nanometers to some tens of

nanometers.

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 relationship between current and voltage, this will no longer be the case in the quantum mechanical regime, as experimental prototypes clearly show.

1.1.6 Energy dissipation in computers




Figure 1.4: Exponential reduction of the energy dissipated per logical operation. The horizontal line labeled " $k_B T$ " indicates the thermal energy per degree of freedom at room temperature.

Possibly even more impressive than the reduction in feature size over time is a corresponding trend in the energy dissipated in a logical step, which is represented in Figure 1.4. Over the last fifty years, this number has decreased by more than ten orders of magnitude, again following an exponential time dependence. The reduction of energy dissipation is a requirement for a continuing improvement: if today's processors had the energy efficiency of 1950, even a single computer would require more power than a large power station can generate. As a result, it would become too hot to operate and disintegrate within a fraction of a second. Even with today's efficiency,

heat dissipation is often the limiting factor for the speed of microprocessors.

While the continued reduction of energy dissipation is thus a necessity, the continuation of this trend is approaching a fundamental limit: Any computer working with Boolean logic (which includes all of today's digital computers) must dissipate at least an amount of $k_B T \ln 2$ whenever it performs an AND or OR operation, since these operations are not reversible.

The $k_B T$ limit could be overcome by using so-called reversible logical operations. As we will discuss in section 5, devices that operate by the laws of quantum mechanics are inherently reversible. The principles on which it operates may thus well find applications also in electronic circuits for classical computers. The two trends – reduction of dissipated power and reduction of size – therefore appear to converge towards devices that use quantum mechanics for their operation.

While the limitations that force the use of quantum devices in the future may appear as a nuisance to many engineers, they also represent an enormous potential, since these future devices may be much more powerful than conventional (classical) devices. They can implement all the algorithms that run on today's classical computers, but in addition, they also can be used to implement a different class of algorithms, which explicitly use the quantum mechanical nature of the device. A few such quantum algorithms have been designed to solve specific problems that cannot be solved efficiently on classical computers. ² While many questions remain unanswered concerning the feasibility of building devices that fulfill all the stringent requirements for a useful quantum computer, the possibilities offered by this emerging technology have generated a lot of attention, even outside the scientific community.

²As discussed in section 3, the term "efficient" means in computer science that the computational cost increases at most polynomially with the size of the problem.

1.2 Quantum computer basics

1.2.1 Quantum information

We discuss here exclusively digital representations of information. Classically, information is then encoded in a sequence of bits, i.e., entities that can be in two distinguishable states, which are conventionally labeled with 0 and 1. The physical representation can assume various forms like black / white color on a QR code, magnetization direction on a hard disk or current of a transistor. In many electronic devices, these states are encoded by voltages, whose values vary with the technological basis of the implementation (e.g.. TTL: 0 ~ low is represented by voltages < 0.8 V and 1 ~ high by voltages > 2.4 V).




Figure 1.5: Representation of information in a classical computer (left) vs. quantum computer (center). The spin 1/2 (right) is the prototypical example of a qubit.

The same principle applies to quantum systems that represent information: to represent a single bit of information, two distinguishable states of the system are needed. "Distinguishable" means, in a quantum system, that the two states must differ in some quantum numbers, i.e., they must be different eigenstates of at least one operator. 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), as shown in the center of figure 1.5.

The main difference between quantum mechani-

cal 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. To emphasize this difference between quantum and classical bits, the term “qubit” (short for quantum bit) has been adopted for the quantum mechanical unit of information.

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. A system consisting of N qubits has 2^N mutually orthogonal basis states. Due to the rapid growth of the exponential function, this results in a very large number of basis states: even for $N = 300$, the number of distinct states exceeds the number of atoms in the universe.

The power of quantum computing is closely related to the possibility of bringing such a system into a state that is a superposition of all basis states. Logical operations such as multiplications can then be applied to this superposition. In a sense to be discussed later, such a transformation is equivalent to transforming all the states in parallel, i.e., performing 2^N operations in parallel.

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 $|\psi\rangle = a|\psi_0\rangle + b|\psi_1\rangle$. 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 qubit can store an infinite amount of information. To obtain the information content, one has to take the measurement process, which retrieves the information, into account: it is never possible to measure exactly the quantum state of a single photon. A single measurement (more precisely: an ideal quantum mechanical measurement as

postulated by von Neuman) 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 qubit 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 will be discussed in section 4.2.11, 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.

1.2.2 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 its 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, a photon qubit can transmit up to two classical bits of information. 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 eavesdropper intercepts a photon, thus detecting that information is being transferred, and that he subsequently re-emits a similar photon to the original receiver, he cannot send an exact copy of the original photon. This necessarily allows the two partners who are trying to establish a secure communication to realize that their communication is being monitored – not for individual photons, but from a statistical analysis of the successfully transmitted photons.

This is not automatic, however. If the communication protocol were to use only the presence or absence of the photon as the information, the eavesdropper would be able to use QND (=quantum nondemolition detection) to observe the passage of the photon. Such experimental schemes [2] can measure a given quantum mechanical variable (such as the 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.

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, switching the axes of their polarizers randomly between two predetermined positions. If the photon pairs are originally in a singlet state, each partner knows then the result of the other partner's measurements provided that 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 (but not the results of their measurements). They can then eliminate those measurements where only one partner detected a photon as well as those for which the orientation of their polariz-

ers were different. Assuming an ideal system, the remaining measurement results are then exactly anti-correlated. If an eavesdropper (usually called Eve) tried to listen in on their communication, her measurements would inevitably affect the transmitted data. A statistical analysis of the measurement results obtained by Alice and Bob, for which they publicly exchange a fraction of their bits, would then reveal the presence of the eavesdropper. This scheme has been tested successfully in a number of experiments by using optical fibers or beams through free space.

China and specifically USTC have been at the forefront of this research. The milestones that were reached include the Micius satellite for satellite-based quantum communication, the quantum-secure connection between Shanghai, Hefei and Beijing, and the metropolitan quantum network in Hefei.

1.2.3 Basics of quantum information processing

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 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 state of the quantum register is written as

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

While the quantum registers in current devices have only limited control over a limited number of qubits, useful quantum computers require several thousand, possibly millions of qubits, depending on the application.

Before an actual computation can be initiated, the quantum register must be initialized into a

well defined state, typically the quantum mechanical ground state $|0, 0, \dots, 0\rangle$. This operation is non-unitary, since it must bring the system into one specific state, independent of the state in which it starts. The initialization is therefore a non-reversible process that must include dissipation.

The actual information processing occurs through the operation of quantum gates, i.e., unitary transformations \mathbf{U}_i that operate on the quantum register and correspond to logical operations:

$$|\psi_0\rangle \xrightarrow{\mathbf{U}_1} |\psi_1\rangle \xrightarrow{\mathbf{U}_2} |\psi_2\rangle \dots$$

The sequence of quantum gates is determined by the specific algorithm to be implemented. The program that specifies this sequence may be stored in a classical device associated with the quantum computer, such as a classical computer.

Like any change in a quantum mechanical system, logical operations are driven by a suitable Hamiltonian acting on the state that represents the quantum register. It is in most cases difficult to find a Hamiltonian that directly performs the desired transformation, such as the decomposition of an integer into its prime factors. Instead, the total transformation is usually split into elementary logical operations that transform a single bit of information or connect two bits by operating on one bit in a way that depends on the state of the other bit. It turns out that every possible logical operation can be constructed by concatenating elementary gate operations that belong to one of two groups:

- single qubit operations, corresponding to arbitrary rotations of the spinor representing the qubit and
- one type of 2-qubit operations, e.g., the “controlled NOT” or CNOT.

A quantum computer implementation that can perform arbitrary calculations must therefore implement these two types of operations. Particularly critical are the two-qubit operations, since they require interactions between the qubits. A

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

Table 1.1: Truth table of CNOT gate.

typical operation is the CNOT gate, whose truth table is shown in table 1.1: this particular gate has two inputs and two outputs. If the control qubit is zero, it simply passes both qubits to the output. If the control qubit is one, it passes the control qubit through unchanged, but inverts the target qubit.

It must be possible to apply the 2-qubit operations to any pair of qubits, as the algorithm demands. This requires in principle the existence of a physical interaction between the two qubits. If such an interaction is not present, it is possible, however, to decompose the corresponding 2-qubit operation into a series of 2-qubit operations between nearest neighbours. Such schemes that use interactions only between nearest neighbors, 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.

1.2.4 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, quantum computers derive their power from the possibility of performing 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. As the

number of qubits in the quantum register increases, the processing power increases, but at the same time the quantum information becomes more fragile.

The biggest contribution to decoherence is usually dephasing. In a simple picture, dephasing occurs when the energy difference between the two states representing the qubit fluctuates. As a result, the relative phase of the superposition state acquires an additional phase proportional to the energy change.

The effect of such a dephasing as well as other decoherence processes 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 processing algorithm before decoherence becomes noticeable, 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 corrections. Basically these processes use coding of quantum information in additional qubits. Algorithms have been developed for using these additional qubits to check for and eliminate various types of errors. Basics of decoherence, errors and quantum error correction will be discussed in chapter 7.

1.2.5 The network model

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 all implementations must fulfill and discuss some issues that help in identifying suitable systems. More details will be discussed in chapters 9-12.

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 as large as possible; limitations on the number N of identifiable qubits will therefore be an important consideration.




Figure 1.6: Principle of operation of quantum processors using the so-called network model.

Figure 1.6 shows the principle of operation of quantum processors using the so-called network model. 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 $U_i = e^{-iH_i\tau_i}$ 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. 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 an eigenstate of the corresponding observable.

1.2.6 Physical implementation

DiVincenzo's criteria

A quantum mechanical system that can be used as an information processing device must meet a number of rather restrictive conditions, including:^[3]

- It must be possible to initialize the system

into a well-defined quantum state.

- It must be possible to apply unitary operations to each individual two-level system that serves as a qubit.
- It must be possible to apply unitary operations to some pairs of qubits.
- The information stored in the quantum register, in particular the relative phases of all quantum states must be preserved for a sufficiently large number of logical operations.
- It must be possible to read out the state of each qubit with high fidelity.

Each of these requirements can be expressed much more quantitatively, as we discuss later in this book. Some of the requirements tend to work against each other: being able to control individual qubits, e.g., requires coupling them to the environment. Such couplings, however, always tend to bring noise into the system, thus causing decay of the quantum information. The same processes that must be used to initialize the system again tend to destroy the quantum information. This will be discussed in more detail in chapter 9 ff.

Satisfying all these requirements in a single physical system is a huge challenge, but over the years many systems have been proposed and for some of them, proofs of principle have been demonstrated that it may be possible to overcome these challenges.

NMR

The first physical system on which quantum algorithms were actually implemented was nuclear magnetic resonance (NMR) in liquids [4, 5]. This is also the most direct realization in the sense that it uses spins 1/2 as qubits. Nuclear spins have the advantage that they are very well isolated from their environment, therefore preserving the quantum information for long times (up to several seconds). On the other hand, the weak coupling to the environment makes it very difficult to measure the spin state of individual

nuclear spins. This difficulty can be circumvented in liquid state NMR quantum computers by working with many identical copies of the molecule that represents the quantum register, typically some 10^{20} . All of the initial demonstrations of QIP therefore relied on liquid-state NMR. The advantage of this scheme is a relatively straightforward implementation of gate operations, the main disadvantage is that such “ensemble” quantum computers are difficult to scale to large numbers of qubits.

Trapped atomic ions

Another physical system that is relatively well isolated from its environment is a system of atomic ions stored in electromagnetic traps [6, 7]. Storing information in these systems is less straightforward than NMR, since the number of states accessible to each ion is infinite and the interactions are harder to control with sufficient precision. The main advantage of trapped ions may be that it is relatively easy to read out the result from individual ions. This is currently a very active field of research in academia and industry; most groups use integrated circuits to trap and control the ions.

Others

In addition to NMR and ion traps, several other systems have implemented simple quantum algorithms. This includes electron spins in solids, photons, trapped neutral atoms, superconducting circuits and quantum dots in semiconductors. Figure 1.7 lists some of them, with attempts to quantify their characteristics.

Solid-state implementations may be easier to scale to larger numbers of qubits, but they are subject to faster decoherence processes, and it remains difficult to control the manufacturing of small structures in a reproducible way. Different companies like IBM quantum computing have realized systems with several hundred qubits.

Which quantum computer is right for you?

There are many types to choose from. Here's how they compare and our all-important verdict




Figure 1.7: Overview of some concepts for quantum computers. (New Scientist, 15 October 2014).

1.3 History of quantum information processing

1.3.1 Initial ideas

Quantum information processing has deep roots that are almost as old as quantum mechanics itself. If we believe that quantum mechanics is the fundamental physical theory that lets us derive properties of all materials, it should also be the basis for the description of any computer. However, in most cases, classical mechanics (and optics, electrodynamics etc.) are excellent approximations to the underlying quantum theory and perfectly adequate for the description of the operation of computational machinery.

The more relevant question is therefore, what happens when the physical basis for the computer is an explicitly quantum system for whose description the classical approximation fails. Explicit discussions on this possibility essentially started in 1980, when Benioff showed how the time dependence of quantum systems could be used to efficiently simulate classical computers operating according to Boolean logic [8].

In the same year, Richard Feynman asked the opposite question: Can classical computers ef-

ficiently simulate quantum mechanical systems [9]. 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 $2 \cdot 2^N - 1$ real numbers. Any computer trying to simulate the evolution of such a system therefore must keep track of 2^N complex numbers. Even for a few hundred particles, 2^N exceeds the number of atoms in the universe and therefore the memory of any conceivable computer that stores these variables in bit sequences. At the same time, the time required to run a simulation grows 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 offered a solution: “Quantum computers – universal quantum simulators”. He showed that the drastic increase in the storage requirements and the computation time can be viewed as a consequence of the large amount of information that is present in the quantum mechanical system. The consideration that quantum systems effectively simulate themselves may then be taken as an indication that they are efficient processors of information. He stated “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 asked which systems could actually be simulated and where such simulations would be useful.

A first proof of this conjecture was given in 1993 by Bernstein and Vazirani [10]. They showed that a quantum mechanical Turing machine is

capable of simulating other quantum mechanical systems in polynomial time. This implied that quantum computers are more powerful than classical computers. This was a proof of principle, but no example was given for such a procedure, i.e., no algorithm was yet known that would run more efficiently on a quantum computer than on a classical computer.

1.3.2 Quantum algorithms

Such algorithms, which require a quantum computer, are called “quantum algorithms”. The first quantum algorithm that can run faster on a quantum computer than on any classical computer was put forward by Deutsch in 1985 [11] and generalized by Deutsch and Jozsa in 1992 [12]. The problem they solved – deciding if all possible results of a function are either identical or equally distributed between two values – had little practical relevance.

A very useful algorithm was developed in 1994 by Don Coppersmith [13]: he showed how the Fourier transform can be implemented efficiently on a quantum computer. The Fourier transform has a wide range of applications in physics and mathematics. In particular it is also used in number theory for factoring large numbers. The best known application of the quantum Fourier transform is the factoring algorithm that Peter Shor published in 1994 [14]. Factoring larger numbers is not only of interest for number theory, but also has significant impact on the security of digital data transmission: The most popular cryptographic systems rely on the difficulty of factoring large 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 [15]. Shor proposed 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 the number of digits. This is a qualitative difference: polynomial-time algorithms are considered “efficient”, while exponential algorithms are

not usable for large systems. The different behavior 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.




Figure 1.8: Time required for classical factorization algorithm vs. quantum algorithm.

We illustrate this by 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, as shown in figure 1.8. 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.

1.3.3 Potential Benefits and Risks

Quantum technologies are relatively recent branch of science and engineering that is evolving rapidly. While there has been very limited impact on areas outside the scientific community, several potential applications are emerging and the number of companies developing quantum computers is increasing rapidly.

Wikipedia (https://en.wikipedia.org/wiki/List_of_companies_involved_in_quantum_computing_or_communication) shows a long and growing lists of such companies.

As the number and range of potential applications is rising, industry and science policy are considering the potential impacts and try to assess them [16]. Benefits are foreseen in areas like finance, healthcare, and materials science. For specific (simple) examples, it has already been demonstrated that quantum algorithms can be superior to classical ones in the prediction of molecular structures. This may result in shorter and less costly development cycles for drugs and materials.

Risks include potential threats to data privacy, national security and cybersecurity, since current public-key cryptographic systems may become vulnerable to attack by powerful quantum computers. Research into safer “post-quantum” cryptographic systems is therefore well under way and standards are being implemented.

Past experience shows that such breakthrough technologies often are difficult to assess and early predictions tend to overestimate the short-term impact but underestimate the long-term effect. This is, to a large degree due to applications that can not be predicted initially, since they do not exist, but emerge as the evolving technology offers qualitatively novel possibilities. Some of the best examples include the laser, which was described as “a solution looking for a problem” when it was invented, or the computer, where its spread was vastly underestimated.

1.4 About this course

1.4.1 Literature

A wide range of textbooks is available on quantum information in general as well as any of its subtopics. A small selection of the general ones is:

- M. A. Nielsen and I. L. Chuang, "Quantum computation and quantum information", Cambridge (2001).
- P. Kaye, R. Laflamme, and M. Mosca, "An introduction to quantum computing", Oxford, (2007).
- D. Bruß and G. Leuchs, "Lectures on Quantum Information", Wiley-VCH (2007).
- M. Nakahara and T. Ohmi, "Quantum Computing: From Linear Algebra to Physical Realizations", CRC press (2008).
- J. Stolze and D. Suter, "Quantum Computing: A Short Course from Theory to Experiment", Wiley-VCH (2008).
- S. Aaronson, "Quantum Computing since Democritus", Cambridge (2013).
- Y. Wan, "Quantum Computing Unveiled", Cambridge (expected 2025).

1.4.2 Online resources

Quantum information processing is an active field of research and reviews of important developments keep appearing in the literature. Perhaps the easiest way to keep in touch with new developments is to consult the “living document” type reports which have been prepared by expert panels in both the USA and Europe and which are (hopefully) updated on a regular basis. The European report “Quantum Information Processing and Communication: strategic report on current status, visions and goals for research in Europe” is available at www.qurope.net or qist.ect.it and its US counterpart “A Quantum Information Science and Technology Roadmap” is available at qist.lanl.gov/qcomp_map.shtml.