

Technical Report No. 2005-496

QUANTUM COMPUTATION AND QUANTUM INFORMATION*

Marius Nagy and Selim G. Akl
School of Computing
Queen's University
Kingston, Ontario K7L 3N6
Canada
E-mail: {marius,akl}@cs.queensu.ca

Abstract

The paper is intended to be a survey of all the important aspects and results that have shaped the field of quantum computation and quantum information. The reader is first familiarized with those features and principles of quantum mechanics providing a more efficient and secure information processing. Their applications to the general theory of information, cryptography, algorithms, computational complexity and error-correction are then discussed. Prospects for building a practical quantum computer are also analyzed.

1 Introduction and overview

Quantum computation and quantum information can be defined as the study of information processing tasks accomplished using quantum mechanical systems. One of the most striking aspects of it is the complete uncertainty about its future. It could rise to meet the expectations of enthusiasts of the field, causing computer science to be reconsidered (and effectively rewritten) in the new *quantum* light. In such a case, quantum computers and practical applications of quantum information theory will substantially impact our everyday life. But it may also happen that quantum mechanics will one day be disproved or the formidable obstacles towards making quantum computers a viable technology will prove insurmountable. This will leave quantum computation and quantum information as abstract mathematical curiosities, without substance.

1.1 Origins of quantum computing

Most people involved in the field associate the birth of quantum computation and quantum information with a talk Richard Feynman gave at MIT in 1981 (see [80] for the journal version). In his talk he pointed out the difficulty of simulating quantum systems using classical computers. In part, this is due to the number of variables a computer must keep track of, when simulating the behavior of a quantum system. This number grows exponentially with the size of the system being

*This research was supported by the Natural Sciences and Engineering Research Council of Canada.

modeled. Other reasons include strange effects that are specific to quantum mechanics and cannot be accurately simulated through classical means. Consequently, Feynman conjectured that a machine built in such a way as to allow it to make use of quantum effects would be able to efficiently simulate quantum systems. And so, the idea of a *quantum computer* was born.

Although Feynman's motivation was the concept of a "universal quantum simulator", a machine capable of imitating any quantum system (including the physical world), from a computer science perspective his observations led to speculation that perhaps computation in general could be done more efficiently if it made use of these quantum effects. In a later paper [81], Feynman analyzes the physical limitations of computers due to the laws of physics and tries to exhibit the principles on which a quantum mechanical computer could be built. As researchers began to explore ways to harness quantum effects in order to speed up computation or find useful applications, the field of quantum computation and quantum information gradually came into its own. In their effort to develop tools that would allow them to gain a better understanding of quantum phenomena, physicists have also contributed to the progress of the field.

Quantum computation and quantum information is a multi-disciplinary area, so it comes as no surprise that people involved in it may have fairly different backgrounds (mathematicians, statisticians, computer scientists, physicists, cryptographers, information theorists, etc.). The power of this novel computational paradigm comes from its foundation: quantum mechanics. Quantum mechanics, physicists say, is the most complete and accurate description of the world we currently have. It is a mathematical theory consisting of rules and principles defining a framework used to develop physical theories. In this sense (as a mathematical theory) quantum mechanics is not difficult to understand. Familiarity with the tools linear algebra provides for manipulating state vectors in Hilbert spaces is all a computer scientist needs in order to tackle quantum information processing tasks. The difficulty appears when we try to apply quantum mechanics to understanding the structure of some complicated molecules or grasping the nature of the forces responsible for keeping the particles in atomic nuclei together.

1.2 Interpretations of quantum mechanics

From a broader perspective, the difficulties associated with quantum mechanics lie in finding a correspondence between its predictions and what happens in the physical reality. In other words, it is a matter of the *interpretation* given to quantum mechanics. It is in this context that we must understand Neils Bohr's statement that "anyone who thinks he understands quantum theory doesn't really understand it". This sentence expresses the fact that we cannot use our everyday life experience and common sense intuitions to gain insights into quantum phenomena.

1.2.1 The Copenhagen interpretation

Neils Bohr, one of the founders of quantum theory and apparently a very determined and strongly willed person, is credited to have imposed what we now call the *Copenhagen interpretation*, which became the standard view among many physicists. According to this view, it doesn't make sense to ascribe intrinsic properties (such as position or velocity) to isolated quantum entities (such as electrons, photons or other elementary particles). The properties of quantum systems only make sense in light of the measurements we make. Taken to its limit, the Copenhagen interpretation denies the reality of an individual photon or electron until it has been observed.

Let us see how different interpretations of quantum mechanics can explain the result of the *double-slit experiment*. It was first conducted by Thomas Young in 1801 and demonstrated that light behaves like waves. In his experiment, Young projected light onto a screen through a barrier




Figure 1: Young's double-slit experiment.

pierced with two closely spaced slits (see Figure 1). What he observed on the screen was an *interference* pattern, the hallmark of waves. The importance of modern-day versions of Young's experiment is best illustrated by Richard Feynman in his *Lectures* [82]. He believed that the result of the double-slit experiment was the fundamental mystery of quantum mechanics.

If Young performed his experiment using simple screens and candlelight, the tremendous advances in technology allow us today to repeat the experiment with very weak light, that is, light produced as one photon at a time. Thus, it is very unlikely that several photons would be found within the experimental apparatus at the same time. Surprisingly (and against our intuitions), given that enough time elapses as to allow the photons, arriving one at a time, to accumulate on the screen, the same interference pattern will appear. The obvious question is: what was each photon interfering with, if it was alone in the experimental apparatus?

According to the Copenhagen interpretation, the only possible answer can be: with itself. In the absence of any observations, it doesn't make sense to associate a specific route to the photon in its way from the light source to the screen. In a sense, each particle went not through one slit, but rather through both slits, and as it appeared on the other side, it interfered with itself. As we will see in the next section, this behavior is a manifestation of the quantum principle of superposition of states, a principle without which quantum computation and quantum information would be unconceivable. The duality between particles and waves has also been demonstrated for other quanta that can be localized (electrons, neutrons, atoms) and even for larger entities, like complex molecules composed of tens of atoms.

1.2.2 Many worlds interpretation

Although the Copenhagen interpretation is over half a century old, there still are physicists today who embrace it. Suffices to name Anton Zeilinger, who was involved in the experimental demonstration of superdense coding and quantum teleportation, two of the most important applications of quantum information theory. However, tentatives of applying quantum mechanics to the universe as a whole led to renewed interest in the many-universes interpretation, initially proposed by Hugh Everett III in 1957. His interpretation essentially differs from the Copenhagen interpretation be-

cause it removes completely the measurement problem, by eliminating any references to observers. According to the modern version of the many-worlds interpretation, our universe is embedded into an infinitely larger and more complex structure, called the *multiverse*, which we can approximate to a system of parallel universes [64, 63, 42]. Each time there is a decision at the quantum level, such as the probability of a radioactive atom to decay or a photon passing through a polarizing filter, the assemblage of universes differentiates along different paths.

How does this interpretation explain the interference pattern created by single particles? Since there is a choice regarding which route the particle will take, at that point the universe will split into two distinct universes. In one of them, the particle goes through the first slit, while its copy from the second universe will traverse the barrier through the second slit. On the other side of the barrier, the two universes will interfere with each other creating a series of fringes on the screen. So, according to the many-universes interpretation, the particle is not interfering with itself, but rather with its copy from a parallel universe. Although it sounds like a rather crazy idea, the multiverse structure offers sometimes a more intuitive explanation to some quantum phenomena than the classical Copenhagen interpretation. David Deutsch, who may be regarded as the strongest advocate of the many-worlds interpretation (and also had some remarkable contributions to the development of quantum complexity theory), even proposed an experiment that would test the existence of multiple universes and therefore be able to distinguish between the different interpretations of quantum theory [60].

1.3 Overview

The remaining sections of this paper try to address the fundamental concepts and questions underlying quantum computation and quantum information. Is indeed this novel approach to performing computation more powerful than the one in use today? How much more powerful? What are those information processing tasks that prove its superiority? Can quantum effects only speed up computation or are there tasks that can only be performed through quantum means? Can we identify the conceptual elements responsible for the power exhibited by a quantum computer? And, finally, is it possible to build such a machine that would exploit quantum effects in order to achieve unprecedented computational power?

The next section describes the fundamental principles of quantum mechanics on which quantum computation and quantum information is based. Section 3 is constructed in analogy with classical information theory and discusses two simple (yet important) quantum information processing tasks: superdense coding and quantum teleportation. Section 4 is concerned with the most promising practical application of quantum information (at least in the short term): quantum cryptography. The most important quantum algorithms conceived up to now are presented in Section 5. Section 6 introduces the main quantum complexity classes and their relations with classical ones. Any computing technology, regardless of the physical implementation, is subject to various types of errors, and quantum computing makes no exception. Section 7 analyzes possible solutions to this problem. Section 8 is a review of physical realizations of quantum bits and logical operations performed on them (quantum gates). Conclusions and final remarks are presented in Section 9.

2 Fundamentals of quantum information and quantum computation

The field of quantum information and quantum computation is based on the postulates governing quantum mechanics. The aim of this section is to provide a description of these postulates and

mathematical formalisms required to work with them to the extent needed for quantum computing. Good introductions to quantum mechanics for computing scientists can be found in [139, 155, 146, 98, 27], but for detailed expositions of the field one should see [82].

Quantum mechanics takes place in the framework provided by linear algebra. We can associate to any isolated physical system a complex vector space with an inner product defined on it, known as the state space of the system. Mathematically, such a vector space with an inner product is called a Hilbert space. At any given point in time, the system is completely described by its state vector, which must be a unit vector in the system's state space.

Quantum state spaces and the transformations acting on them are traditionally described in terms of vectors and matrices using the compact *bra/ket* notation introduced by Dirac [70]. According to his conventional notation, *kets* like $|x\rangle$ are simply column vectors, typically used to describe quantum states. Similarly, the matching *bra* $\langle x|$ is a row vector denoting the conjugate transpose of $|x\rangle$.

2.1 The qubit

At an abstract level, the simplest quantum mechanical system is the quantum bit, or *qubit*. A qubit is a unit vector in a two-dimensional state space, for which a particular orthonormal basis, denoted by $\{|0\rangle, |1\rangle\}$ has been fixed. The two basis vectors $|0\rangle$ and $|1\rangle$ correspond to the possible values a classical bit can take. However, unlike classical bits, a qubit can also take many other values. In general, an arbitrary qubit $|\Psi\rangle$ can be written as a linear combination of the computational basis states:

$$|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{1}$$

where α and β are complex numbers such that $|\alpha|^2 + |\beta|^2 = 1$. This is the fundamental difference distinguishing quantum bits from classical ones and is a direct application of the quantum principle of superposition of states. The qubit $|\Psi\rangle$ in equation 1 is in a superposition of $|0\rangle$ and $|1\rangle$, a state in which it is not possible to say that the qubit is definitely in the state $|0\rangle$, or definitely in the state $|1\rangle$. After all, what better intuition about the superposition principle than the idea (quite old and widely accepted now) that each particle is also a wave.

For a single qubit, there is a very intuitive graphical representation of its state as a point on the unit three-dimensional sphere (also called the *Bloch sphere*). Figure 2 depicts four possible states of a qubit using the Bloch sphere representation. Note that the states corresponding to the points on the equatorial circle have all equal contributions of 0-ness and 1-ness. What distinguishes them is the *phase*. For example, the two states displayed above, $1/\sqrt{2}(|0\rangle + |1\rangle)$ and $1/\sqrt{2}(|0\rangle - |1\rangle)$ are the same up to a relative phase shift of π , because the $|0\rangle$ amplitudes are identical and the $|1\rangle$ amplitudes differ only by a relative phase factor of $e^{i\pi} = -1$.

We have described qubits as mathematical objects, but there are real physical systems which may be described in terms of qubits. Possible physical realizations of a qubit include two different polarizations of a photon, the alignment of a nuclear spin in a uniform magnetic field or two electronic levels in an atom. There will be more to say about this in section 8.

2.2 Measurements

We now turn our attention on the amount of information that can be stored in a qubit and, respectively, retrieved from a qubit. Since any point on the Bloch sphere can be characterized by a pair of real-valued parameters taking continuous values, it follows that, theoretically, a qubit could




Figure 2: The Bloch sphere representation of a qubit.

hold an infinite amount of information. However, we cannot extract more information from such a qubit than we are able to do it from a classical bit. The reason is that we have to *measure* the qubit in order to determine which state it is in. And another of the fundamental postulates of quantum mechanics, the one regarding measurements, restricts us in the amount of information that can be gained about a quantum state through measurement. According to this postulate, when we measure a qubit $|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle$ with respect to the standard basis for quantum computation $\{|0\rangle, |1\rangle\}$, we get either the result 0 with probability $|\alpha|^2$, or the result 1 with probability $|\beta|^2$. The condition that the probabilities must sum to one corresponds geometrically to the requirement that the qubit state be normalized to length 1, that is the inner product $\langle\Psi|\Psi\rangle$ equals 1.

Furthermore, measurement alters the state of a qubit, collapsing it from its superposition of $|0\rangle$ and $|1\rangle$ to the specific state consistent with the measurement result. For example, if we observe $|\Psi\rangle$ to be in state $|0\rangle$ through measurement, then the post-measurement state of the qubit will be $|0\rangle$, and any subsequent measurements (in the same basis) will yield 0 with probability 1. In general, measurement of a state transforms the state into one of the measuring device's associated basis vectors. The probability that the state is measured as basis vector $|u\rangle$ is the square of the norm of the amplitude of the component of the original state in the direction of the basis vector $|u\rangle$.

2.3 Uncertainty principle

As noted by Julian Brown (in [42] page 120), Berthiaume and Brassard connect Heisenberg's uncertainty principle with the idea that any measurement yielding information in the classical form induces an irrevocable destruction of some remaining information, making further measurements less informative, if not completely useless [28]. Our objection to this point of view is that while it is perfectly true that, in general, measurement disturbs the quantum state of the observed system, this is not what Heisenberg's uncertainty principle is trying to quantify. Heisenberg showed that it is not possible, according to quantum mechanics, to simultaneously know both the position and momentum of a particle with arbitrary accuracy. Position and momentum form a pair of canonically

conjugate quantum variables and the more we know about one of them, the more we are restricted in the precision with which the other one can be measured.

Later, Heisenberg's uncertainty principle was generalized to any two quantities corresponding to non-commuting operators [157]. An equivalent formulation of this general uncertainty relation is developed in [146] based on expressing projective measurement and standard deviations in terms of observables. Of course, there are many other pairs of conjugate variables (beside position and momentum), and we only mention here energy and time (who actually needs a special treatment, as given in [167]), horizontal and vertical polarizations or right-handed and left-handed circular polarizations. It is also important to stress that uncertainty relations of any kind do not express the incapacity of the measurement device to provide accurate observations of the variables, due to their imperfections. These relations are mathematical results derived from the principles governing quantum mechanics and therefore hold regardless of the performances of the measuring devices used.

2.4 No-clonability

Naturally, measurements in bases other than the computational basis are always possible, but this will not help us in determining α and β from a single measurement. One might think of solving this problem by making multiple copies of the initial qubit $|\Psi\rangle$ and then measure each of the copies in order to obtain an estimation of α and β . In fact, it turns out to be impossible to make a copy of an unknown quantum state. The *no-cloning* theorem, one of the earliest results of quantum computation and quantum information [184], states that quantum mechanics prevents us from building a quantum cloning device capable of copying non-orthogonal quantum states. The ability to clone orthogonal quantum states translates into the ability to copy classical information, since the different states of classical information can be thought of merely as orthogonal quantum states. So it seems that quantum mechanics places severe limitations on the accessibility of quantum information, but sometimes this can be used to our advantage, as we will see in the case of quantum cryptography.

2.5 Superposition and interference

The concepts of superposition, interference and measurement can be very well illustrated in the case of Young's two-slit experiment. The experimental setup provides the particle with a particular kind of superposition. If we ascribe state $|0\rangle$ to the particle when it passes through slit A and state $|1\rangle$ when it passes through slit B, then, effectively, the particle's behavior can be described by the superposition of states $\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$, which is a combination of "particle goes through slit A" with "particle goes through slit B". In the case when we choose to observe the particle as it goes through the experimental apparatus (that is, to measure its state), the wave function describing it will collapse into one of the two possible outcomes and the particle will be detected passing through slit A or B with equal probability. In either case, the superposition is destroyed and with it any chance of interference. But if the particle is not observed until the end, as it collects on the screen, then the superposition holds through to the end, enabling the interference phenomenon witnessed on the screen.

2.6 Quantum registers

Let us examine now more complex quantum systems, composed of multiple qubits. In classical physics, individual two-dimensional state spaces of n particles combine through the cartesian prod-

uct to form a vector space of $2n$ dimensions, representing the state space of the ensemble of n particles. However, this is not how a quantum system can be described in terms of its components. Quantum states combine through the tensor product to give a resulting state space of 2^n dimensions, for a system of n qubits. It is this exponential growth of the state space with the number of particles that quantum computers try to exploit in their attempt to achieve exponential speed-up of computation over classical computers.

For a system of two qubits, each with basis $\{|0\rangle, |1\rangle\}$, the resulting state space is the set of normalized vectors in the four dimensional space spanned by basis vectors $\{|0\rangle \otimes |0\rangle, |0\rangle \otimes |1\rangle, |1\rangle \otimes |0\rangle, |1\rangle \otimes |1\rangle\}$, where $|x\rangle \otimes |y\rangle$ denotes the tensor product between column vectors $|x\rangle$ and $|y\rangle$. It is customary to write the basis in the more compact notation $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$. This generalizes in the obvious way to an n -qubit system with 2^n basis vectors.

2.7 Entanglement

Similar to single qubits, multiple-qubit systems can also be in a superposition state. The vector

$$|\Psi\rangle = \frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle) \quad (2)$$

describes a superposition state of a two-qubit system in which all four components (corresponding to the four basis vectors) have equal amplitudes. What about the two qubits composing the system? Can we characterize their states individually? If we rewrite equation 2 in order to express $|\Psi\rangle$ as the tensor product

$$|\Psi\rangle = \left(\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle\right) \otimes \left(\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle\right) \quad (3)$$

then we can legitimately assert that each of the component qubits is also in a superposition state, perfectly balanced between $|0\rangle$ and $|1\rangle$. Now let us drop the two middle terms in equation 2 and consider the superposition state described by

$$|\Phi\rangle = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle \quad (4)$$

In this case it is no longer possible to find complex numbers α , β , γ and δ such that

$$(\alpha|0\rangle + \beta|1\rangle) \otimes (\gamma|0\rangle + \delta|1\rangle) = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle \quad (5)$$

The state of the system cannot be decomposed into a product of the states of the constituents. Even though the state of the system is well defined (through the state vector $|\Phi\rangle$), neither of the two component qubits is in a well-defined state. This is again in contrast to classical systems, whose states can always be broken down into the individual states of their components. Furthermore, if we try to measure the two qubits, the superposition will collapse into one of the two basis vectors contributing to the superposition and the outcomes of the two measurements will always coincide. In other words, if one of the qubits is found to be in state $|0\rangle$, then the second one will necessarily be in the same state, while a state $|1\rangle$ assumed after measurement will be shared by both qubits. Therefore, we say that the two qubits are entangled and $|\Phi\rangle$ describes an entangled state of the system.

Entanglement defines the strong correlations exhibited by two or more particles when they are measured and which cannot be explained by classical means. This does not imply that entangled particles will always be observed in the same state, as entangled states like

$$\frac{1}{\sqrt{2}}|01\rangle \pm \frac{1}{\sqrt{2}}|10\rangle \tag{6}$$

prove it. States like these or the one in equation 4 are known as Bell states or EPR pairs after some of the people (see below) who pointed out their strange properties.

In some sense, we can say that superposition encompasses entanglement, since entanglement can be viewed as a special case of superposition. It is also interesting to make an analogy between entanglement and the concept of primality from number theory. Indeed, an entangled state of the system corresponds to a prime number, since it cannot be factored or decomposed as a product of subsystem states.

2.7.1 EPR thought experiment

Entanglement is by far the most counterintuitive concept in quantum mechanics that we have met. It can give rise to strange phenomena that have fueled endless debates between the advocates of the theory of quantum mechanics and those trying to disprove it (the arguments exchanged by Bohr and Einstein on this subject have become history [2]). Einstein was deeply dissatisfied with the fact that quantum mechanics allowed correlations between entangled particles to manifest themselves instantaneously over arbitrary large distances.

According to quantum mechanics, if the two particles in an EPR pair are arbitrarily far apart and we measure one of them, the other will instantly acquire a definite state. This seemingly instant influence over possibly infinite distances, unmediated by any communication or propagation medium is called *non-locality*, but Einstein termed it “spooky action at a distance” ([2] page 122). He believed that such phenomenon is due to some hidden variable in the state of the system that we simply do not have knowledge of. Supporters of this local hidden variable theory assume that each particle has some internal state (for the moment hidden from us) that completely determines what the result of any given measurement will be. The strong correlation exhibited, for example, by the EPR pair in equation 4 has in their view a simple explanation: the particles are either both in state $|0\rangle$ or both in state $|1\rangle$, we just don’t happen to know which.

Einstein presented his viewpoint in a 1935 paper ([74]) written together with Podolsky and Rosen (hence, the EPR acronym). They imagined a thought experiment that would either show quantum mechanics as being an incomplete theory of Nature or, basically, contradict Heisenberg’s uncertainty principle. Bohr’s reply came only a few months later ([33]) pointing out the flaw in Einstein’s argument: two entangled particles form an inseparable quantum system and any measurement carried out on one of them will inevitably influence the circumstances of the other. But the dispute between adepts of realistic local theories and supporters of quantum mechanics (and implicitly entanglement) was far from being over.

2.7.2 Bell’s inequality

John Bell set up a new battlefield in 1964 ([12]), in the form of a mathematical expression that clearly separates classical from quantum correlations. This expression, which takes the form of an inequality known as *Bell’s inequality* quantifies the limit on the strength of correlations we can expect to see if hidden variables are involved. On the other hand, quantum mechanics ought to violate this inequality. Bell’s inequality was only the first in a larger set of inequalities of this kind. Another example is the CHSH inequality, named after the initials of its four discoverers [53].

Bell’s result prompted a series of experiments arduously aimed at demonstrating that the laws of Nature do indeed violate Bell’s inequality. In the most famous of them, the correlations in polar-

izations between entangled photons was measured [3]. Alain Aspect and his colleagues claimed that the results confirmed the violation of Bell’s inequality, thus proving the existence of entanglement in Nature and the validity of quantum mechanics. Later experiments demonstrated quantum-correlations over larger distances [170, 175]. But there are voices who doubt the results of these experiments, accusing a lack of justification for all the assumptions on which the experiments are based [134, 172, 173, 174]. They consider the attitude of the experimenters to be biased and the whole demonstration of quantum entanglement as a circular argument: the assumptions needed to back quantum mechanics are unlikely to be true unless quantum mechanics is true.

Any experiment carried out to prove that quantum mechanics violates the Bell inequalities requires an extended run of measurements. The frequencies with which pairs of measurements agree with one another are then calculated. It is claimed that the EPR effect has been convincingly demonstrated with statistically significant results. But it is exactly the statistical nature of the argument that has been pointed out as a weakness. Critics of these experiments consider that their statistical results have been biased by resorting to convenient assumptions, insufficiently founded.

2.7.3 3-particle entanglement

Under these circumstances, conceiving a new kind of experiment that would allow the clash between quantum and classical reality to be decided in one measurement seemed an important achievement. The GHZ experiment [92] is basically an extension of the EPR experiment with three correlated particles instead of two. The three-particle entanglement in the GHZ proposal provides the means to prove the contradiction without the cumbersome use of inequalities, in a much more direct and non-statistical way, as compared with Bell’s original theorem.

In 1999 a team of researchers at MIT produced a GHZ state using nuclear spins instead of photon polarizations and the techniques of nuclear magnetic resonance spectroscopy to manipulate their sample [144]. Although they were able to measure the weird quantum correlations exhibited by the GHZ state, the NMR techniques used prevented them from testing the non-local aspects of the GHZ experiment. The three qubits, embodied as nuclear spins within a molecule, were far too close together to allow for a delayed-choice experiment of the kind Aspect performed on EPR pairs. There is also a very recent report about the successful creation of a 3-particle entangled GHZ state between trapped ions [55], as part of the effort to develop a scalable technology for building a quantum computer.

2.7.4 Entanglement as a physical resource

Today, entanglement is no longer regarded as merely a quantum curiosity. It has become the essential feature of quantum mechanics and people are seeing it as a *physical resource* that can be spent in order to solve information-processing tasks in new ways. Some researchers have even begun to quantify entanglement by introducing a standard unit of entanglement in the form of a maximally entangled pair, which is said to carry one *e-bit* of entanglement [145]. By contrast, incompletely entangled pairs carry less than one *e-bit*. They developed ways to weigh entanglement using entanglement distillation and dilution. Therefore, **the question whether this “resource” really exists in Nature and can be exploited in the ways we currently believe it can, is of crucial importance for the future of quantum computation and quantum information**¹. Superdense coding, teleporting a quantum state or finding the prime factors of a composite integer in polynomial time on a quantum computer are only a few of the most important applications of

¹Text in **bold** is meant to highlight important open problems or areas of intense research.

entanglement. The current potential quantum computation and quantum information are credited with will certainly be affected if some day entanglement will be disproved.

2.8 Quantum evolution

Let us assume for now that we can harness entanglement to increase the efficiency of computation and information-processing tasks in general. What are the “circuits” composing a hypothetical quantum computer? The answer to this question is in strong relation with the way an isolated quantum system evolves over time. We already saw what happens when we try to measure such a quantum system. If, for example, we are trying to read the content of a quantum memory register, the system will undergo a sudden, unpredictable jump into one of the classical bit string configurations composing the original superposition. In other words, there will be a discontinuity in the evolution of the quantum memory register. But, if we leave the register unobserved, the system will undergo a smooth, continuous evolution governed by Schrödinger’s equation, a deterministic differential equation which enables us to predict the future or uncover the past evolution of the memory register. Consequently, any quantum computation is reversible and therefore quantum gates (the quantum analog of classical gates) must always have as many outputs as they have inputs, in order to avoid any loss of information that would prevent the computation to be undone.

2.9 Quantum gates

A quantum NOT gate acting on a single qubit will evolve the initial state $\alpha|0\rangle + \beta|1\rangle$ into the final state $\alpha|1\rangle + \beta|0\rangle$, in which the roles of $|0\rangle$ and $|1\rangle$ have been interchanged. Because every quantum gate acts linearly, the transformation is fully specified by its effect on the basis vectors. Hence, there is a very convenient representation of a quantum gate in matrix form. The matrix X representing the quantum NOT gate is then defined as follows:

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

The first column represents the effect of applying the NOT gate to state $|0\rangle$, while the second column is the result of applying the NOT gate to state $|1\rangle$. We can now describe the operation of the quantum NOT gate, acting on an arbitrary qubit state, through the following equation:

$$X \cdot \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \beta \\ \alpha \end{bmatrix}.$$

Other examples of single qubit gates are the Z gate:

$$Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix},$$

which leaves $|0\rangle$ unchanged, but introduces a phase shift by flipping the sign of $|1\rangle$, and the Hadamard gate:

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix},$$

which is one of the most useful quantum gates, because it creates superpositions of $|0\rangle$ and $|1\rangle$.




Figure 3: Controlled-NOT quantum gate.

Although there are an infinite number of single qubit gates, not any two by two matrix is a legitimate representation of a quantum gate. Schrödinger's equation states that the dynamics of a quantum system must take states to states in a way that preserves orthogonality. In other words, the normalization condition $|\alpha|^2 + |\beta|^2 = 1$ for the initial state $\alpha|0\rangle + \beta|1\rangle$ must also be true for the quantum state after the gate has acted. This translates into the requirement that the matrix U describing the single qubit gate be *unitary*, that is $U^* \cdot U = I$, where U^* is the conjugate transpose of U . Single qubit gates can be conveniently visualized as rotations of the arrow representing the qubit state on the surface of the Bloch sphere.

Quantum gates on multiple qubits can also be defined. Figure 3 depicts a controlled-NOT gate, an instance of the more abstract controlled- U gate, where $U = X$. The target bit $|y\rangle$ is flipped if and only if the control bit $|x\rangle$ is set to 1.

Multiple qubit gates must also satisfy the requirement that probability be conserved, so they too must be unitary transformations. Since any unitary matrix is invertible and the inverse is also a unitary matrix, it follows that a quantum gate can always be inverted by another quantum gate. The set of all 1-qubit rotations (gates) together with the controlled-NOT gate is universal for quantum computation. But finite universal sets of gates exist as well. Two researchers working independently have shown that any imaginable quantum computation can be performed by connecting together multiple copies of a certain 2-qubit gate [4, 72]. Such universal quantum gates are similar to the NAND gate in classical computation.

2.10 Density operators

The purpose of this section was to introduce the postulates of quantum mechanics on which quantum computation and quantum information processing is based. These postulates were formulated using the language of state vectors. However, there is an alternate formulation, that is mathematically equivalent to the state vector approach, which proves to be much more convenient in some situations, notably as a tool for the description of individual subsystems of a composite quantum system. This tool, known as the *density operator* or *density matrix* provides a convenient means for characterizing quantum systems whose state is not completely known.

Suppose that a quantum system is in one of a number of states $|\psi_i\rangle$, with respective probabilities p_i . Then, by use of outer products, we can describe the state of the system through the following density operator (matrix):

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|. \quad (7)$$

Because the states $|\psi_i\rangle$ are known exactly (they can be represented by a state vector), they are called *pure states*, as opposed to the state described by ρ , which is a *mixed state*. There is a simple criterion to determine whether a state is pure or mixed, using the *trace* of a matrix. A pure state satisfies $tr(\rho^2) = 1$, while for a mixed state, $tr(\rho^2)$ is always less than 1.

Density operators may be regarded as a more general approach than state vectors. This is clearly seen in the analysis of composite quantum systems, where the use of reduced density operators is virtually indispensable. Here is a significant example. According to the definition, the density matrix for the Bell state $(|00\rangle + |11\rangle)/\sqrt{2}$ is:

$$\rho = \left(\frac{|00\rangle + |11\rangle}{\sqrt{2}} \right) \left(\frac{\langle 00| + \langle 11|}{\sqrt{2}} \right) = \frac{|00\rangle\langle 00| + |11\rangle\langle 00| + |00\rangle\langle 11| + |11\rangle\langle 11|}{2}.$$

If we want a description of the state of the first qubit, we can find its reduced density operator by tracing out the second qubit from the matrix representing the joint system:

$$\rho_1 = \text{tr}_2(\rho) = \frac{|0\rangle\langle 0| + |1\rangle\langle 1|}{2} = \frac{I}{2}.$$

This result tells us that the first qubit is in a mixed state, since

$$\text{tr} \left(\left(\frac{I}{2} \right)^2 \right) = \frac{1}{2} < 1.$$

Naturally, a similar result can be obtained for the second qubit. Notice again the paradox introduced by entanglement. Although the joint system is in a pure state (it is known completely), any of the composing subsystems is in mixed states, apparently suggesting that we do not have complete knowledge about it.

From this example we can see that density operators are able to provide a description for the state of a quantum subsystem, even when no state vector can be associated with that subsystem. Yet, in our opinion, density operators fail to capture the essence of entanglement, since a qubit that is either in state $|0\rangle$ or in state $|1\rangle$ with equal probability will yield the same density matrix

$$\rho = \frac{|0\rangle\langle 0| + |1\rangle\langle 1|}{2},$$

as the entangled qubit above.

3 Quantum information theory

Quantum information theory is the field concerned with the study of elementary information processing tasks achieved using quantum mechanics. In a wide context, quantum information theory seeks to identify those resources, separating the quantum from the classical world, which create new possibilities for processing information. More concrete, and in analogy with the classical field, work on quantum information theory can be characterized by the following fundamental goals:

- identify elementary classes of static resources in quantum mechanics (e.g. bits, qubits, entanglement)
- identify elementary classes of dynamical processes in quantum mechanics (e.g. classical or quantum information transmission, noise in a classical or quantum communications channel)
- quantify the resources required to perform elementary dynamical processes.

Since classical physics arises as a special case of quantum physics, all the static and dynamic elements of classical information theory are also present in quantum information theory. However, the latter is broader in scope, for it also includes additional static and dynamic elements, that are specific to quantum mechanics.

3.1 Classical information through quantum channels

The most fundamental results in classical information theory are Shannon's noiseless channel coding theorem and noisy channel coding theorem [161]. It is natural for quantum information theory to look at these two problems in a more general context. In a first step, only the storage medium is changed, so that classical information is transmitted using quantum states as the vehicle. It turns out that using qubits to compress classical information does not result in any significant saving in the amount of communication required to transmit information over a noiseless channel [146]. In the more realistic case, where the communications channel is affected by noise, the channel's true capacity is quantified exactly by Shannon's noisy channel coding theorem. This result proved difficult to replicate for a quantum channel, due to the huge variety of noise models allowed by the continuous space in which quantum evolution takes place. The Holevo-Schumacher-Westmoreland theorem [103, 160] provides only a lower bound on the capacity of such a channel. **It is still a major open problem of quantum information theory to determine whether or not encoding using entangled states can be used to raise the capacity beyond this lower bound.**

3.2 Quantum information through quantum channels

The analogy with Shannon's coding theorems can be taken further, by considering quantum states themselves as the static resource involved in compression and decompression manipulations. Compressing the output produced by a quantum information source is still possible, but the process may no longer be error-free. The quantum states being produced by the source may be slightly distorted by the compression-decompression procedure. The average distortion introduced by a compression scheme is quantified by a *fidelity* measure, analogous to the probability of doing the decompression correctly. Schumacher's quantum noiseless channel coding theorem [159] quantifies the resources required to perform quantum data compression, with the restriction that it be possible to recover the source with fidelity close to 1, in the limit of large block lengths.

Shannon's noiseless channel coding theorem tells us that the number of bits of information necessary to represent, on average, each use of a classical source of information is given by a function of the source probability distribution, called the *Shannon entropy*. Similarly, Schumacher's theorem introduces a new entropic quantity, namely, the *von Neumann entropy*, as the limit to which a quantum source may be compressed. This new entropy can be interpreted as a generalization of Shannon's entropy to quantum states. The von Neumann entropy agrees with Shannon's entropy if the states produced by the quantum source are orthogonal, but in general it is strictly smaller than Shannon's entropy. This decrease in resources required to compress a quantum source is possible exactly by exploiting the inherent redundancy arising in non-orthogonal quantum states. **A fully satisfactory analogue of Shannon's noisy channel coding theorem for encoding and decoding quantum states traveling through a noisy quantum channel has not yet been found.** The capacity has been established however for some specific channels, like the quantum erasure channel [23].

Entanglement is probably the most bizarre elementary static resource of quantum mechanics. Its properties, which are yet to be well understood, are essentially different from those of the resources most familiar from classical information theory. But these strange properties of entanglement are also responsible for creating novel and surprising possibilities of accomplishing information processing tasks. In the following we give two examples of simple, yet unexpected applications of elementary quantum mechanics to the problem of transmitting information between two parties, conventionally known as *Alice* and *Bob*, who are a long way away from one another.

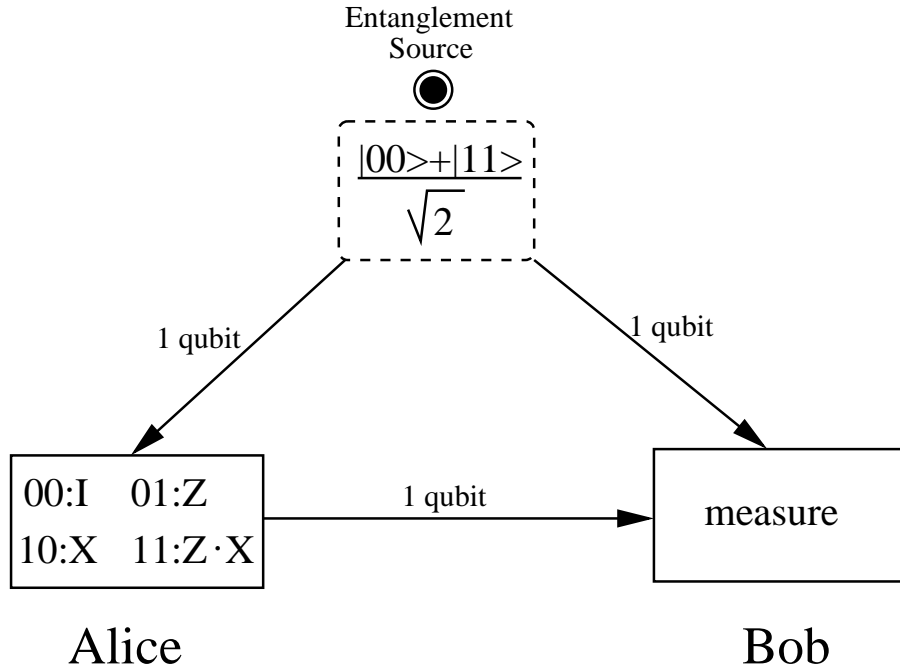


Figure 4: Superdense coding.

3.3 Superdense coding

In the first application, known as *superdense coding* [25], entanglement allows Alice to communicate to Bob two bits of classical information by sending him only one qubit. To do this though, they must previously share a pair of qubits in the entangled state $(|00\rangle + |11\rangle)/\sqrt{2}$. Alice is initially in possession of the first qubit, while Bob has possession of the second one. There are four different messages Alice can transmit over to Bob using two bits of information: '00', '01', '10' or '11'. For each of these four bit strings, Alice applies a specific quantum gate to her qubit and then sends it over to Bob (see Figure 4). Now Bob has hold of both qubits and is free to perform any kind of measurement on them.

Because the two qubits were entangled, any manipulation of the first qubit done by Alice has inevitably affected the state of the ensemble, such that the final state describing the qubits is one of the four Bell states or EPR pairs: $(|00\rangle \pm |11\rangle)/\sqrt{2}$, $(|01\rangle \pm |10\rangle)/\sqrt{2}$. The catch is that the Bell states form an orthonormal basis and can therefore be distinguished by an appropriate quantum measurement. All that Bob has to do in order to determine which of the four possible bit strings Alice sent is to perform a *joint* measurement on the two qubits in the Bell basis.

In terms of the three goals of information theory stated at the beginning of this section, we can identify the two qubits and the entanglement between them as the static resources involved. As dynamical processes we can name the entanglement transformation performed by Alice, the transmission of the qubit from Alice to Bob and the measurement performed by Bob. What are the physical resources that have to be spent in order to complete the task? The answer to this question depends ultimately on the physical realization of the joint measurement. Theoretically, since the state of the 2-qubit ensemble is already perfectly aligned with one of the measurement's projectors, the act of measurement should not change this state. The EPR pair should retain its state throughout the measurement process. As one would expect, such a joint measurement is very difficult to implement in practice. Therefore, assuming that Bob's measurement collapses the

entangled state, the resources consumed in the process are one qubit worth of communication and the entanglement relating the two qubits. Transmitting two bits of classical information through one qubit is only possible by spending one e-bit of entanglement. Of course, the protocol requires two qubits, but Alice never need interact with the second qubit. The initial, entangled state of the ensemble is a fixed state, so there is no need for Alice to have sent Bob any qubits in order to prepare this state. Instead, some third party (an "entanglement provider") may prepare the entangled state ahead of time, sending one of the qubits to Alice and the other to Bob.

3.3.1 Experimental demonstration

The superdense coding protocol has received partial verification in the laboratory. An experiment performed at the University of Innsbruck in Austria [135] implemented a variant of superdense coding using entangled photons and the techniques of *two-particle interferometry*, which involve interference between pairs of particles. The biggest problem was the joint measurement that Bob has to perform to distinguish between the four Bell states. In the Innsbruck experiment, two of the possibilities cannot be distinguished from one another, so Bob can only read three out of four possible states for each photon Alice sends him. Therefore, although this experiment failed to achieve the theoretical two classical bits per qubit, it nevertheless still managed more than one bit (more precisely, a qubit was carrying one *trit* of information).

3.4 Quantum teleportation

Superdense coding is a protocol for transmitting classical information through a quantum channel. We now move to the second example, *quantum teleportation* [20], in which Alice is able to send Bob an unknown quantum state, using only a classical communications channel. Furthermore, the transfer of hidden quantum information implied by the process appears to happen without that state having to move through the intervening space, effectively *teleporting* the quantum state to Bob. Again, this is only possible if Alice and Bob share an EPR pair. So, initially, Alice has in her possession two qubits: the state to be teleported $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where α and β are unknown amplitudes, and her half of the entangled pair $(|00\rangle + |11\rangle)/\sqrt{2}$, while Bob is in possession of the other half. To achieve teleportation, Alice interacts the qubit $|\psi\rangle$ with her half of the EPR pair and then measures the two qubits in her possession, obtaining one of four possible classical results: 00, 01, 10 or 11.

In the first step, Alice applies a controlled-NOT gate to her qubits (with $|\psi\rangle$ acting as the control qubit) changing the initial state of the 3-qubit system

$$|\Psi_0\rangle = \frac{1}{\sqrt{2}}[\alpha|0\rangle(|00\rangle + |11\rangle) + \beta|1\rangle(|00\rangle + |11\rangle)]$$

to

$$|\Psi_1\rangle = \frac{1}{\sqrt{2}}[\alpha|0\rangle(|00\rangle + |11\rangle) + \beta|1\rangle(|10\rangle + |01\rangle)].$$

In the expressions above, the first two qubits belong to Alice and the last one to Bob. The EPR pair consists of the last two qubits.

Notice how this operation has transformed the initial entanglement into another form. No two qubits in $|\Psi_1\rangle$ form a maximally entangled pair, but the entanglement involves now all three qubits. We can say that the initial entanglement has diffused to the first qubit as well.

In the last step before measurement, Alice sends her first qubit through a Hadamard gate, determining the system to evolve into the state

$$|\Psi_2\rangle = \frac{1}{2}[\alpha(|0\rangle + |1\rangle)(|00\rangle + |11\rangle) + \beta(|0\rangle - |1\rangle)(|10\rangle + |01\rangle)].$$

Since $|\Psi_2\rangle$ is a superposition of all possible $2^3 = 8$ classical states for the 3-qubit system, we can intuitively assert that, in some sense, the entanglement dilution process initiated in the first step is now complete. Does this mean that the teleportation was achieved? By regrouping the terms composing the $|\Psi_2\rangle$ state, we can see that Bob's qubit can only be in one of the following four states: $\alpha|0\rangle \pm \beta|1\rangle$, $\alpha|1\rangle \pm \beta|0\rangle$, corresponding to the four possible outcomes for Alice's measurement. Tracing out Alice's system, the reduced density operator of Bob's system can be shown to be

$$\rho^{Bob} = \frac{|0\rangle\langle 0| + |1\rangle\langle 1|}{2} = \frac{I}{2}.$$

This state has no dependence upon the state $|\psi\rangle$ being teleported. Consequently, any measurement Bob might perform on his qubit will contain no information about $|\psi\rangle$. The situation remains unchanged for Bob, even after Alice has completed her measurements². It is only when Bob is informed of the outcome of Alice's measurement that he knows what quantum transformation (rotation) to apply to his qubit in order to recover the state $|\psi\rangle$ accurately, thus completing the teleportation.

Because Alice must transmit her measurement result to Bob over a classical communications channel, which is by definition limited to the speed of light, quantum teleportation cannot be used to transmit quantum states faster than light. The use of the reduced density operator above makes this argument mathematically rigorous. On the other hand, faster than light communication would have generated a series of paradoxes derived from Einstein's theory of relativity (like the possibility to send information backwards in time, for instance).

A superficial glance at quantum teleportation may lead one to believe that the process creates a copy of the quantum state being teleported, thus violating the no-cloning theorem imposed by quantum mechanics. This violation is only apparent though, since after the teleportation process only Bob's qubit is left in the state $|\psi\rangle$, while the original data qubit ends up in one of the computational basis states $|0\rangle$ or $|1\rangle$, following the measurement performed by Alice on the first qubit.

The information-theoretic static resources making quantum teleportation possible are three qubits (two of them entangled in an EPR pair) and two classical bits. The dynamic processes acting on these resources are the various quantum transformations (gates) to which the three qubits are subjected, Alice's measurement of her two qubits and the two classical bits worth of communication between Alice and Bob. As opposed to superdense coding, here the measurements on Alice's two qubits are carried out independently of one another, in the computational basis. Therefore, we can affirm with certainty that teleporting an unknown quantum state consumes one e-bit of entanglement, since there is no entanglement left in the state of the 3-qubit system at the end of the process. From a broader perspective, quantum teleportation emphasizes the interchangeability of different resources in quantum mechanics, showing that one e-bit of entanglement together with two classical bits of communication is a resource at least the equal of one qubit of communication.

²Note how the knowledge about Bob's subsystem (qubit) is relative to the two observers, Alice and Bob.

3.4.1 Experimental demonstrations

Quantum teleportation is not only a neat trick, but has applications of practical importance, like designing quantum gates that are resistant to the effects of noise. Practical verifications of quantum teleportation have been experimentally realized in various different forms by several teams [34, 35, 86, 147, 11]. Zeilinger and his colleagues at Innsbruck [35], for example, used a variation of their dense coding scheme to teleport the polarization state of a photon. It may be worthwhile to also mention the original idea on which Francesco De Martini and colleagues at the University of Rome [34] built their experiment. Instead of measuring two photons, they chose to measure two different aspects of one photon. One aspect was the polarization and the other was the choice between traveling along one of two different routes. As a matter of fact, the entanglement in the initial EPR pair was also in the choice of paths. An interesting novelty was brought by the most recent experiment [11], in which electrons (and not photons) were used to embody the teleported quantum state.

In both examples that we have described in this section, entanglement plays a crucial role in the successful completion of the respective information processing tasks. The importance of entanglement for quantum computation and quantum information has motivated the development of a fascinating area of study about the creation and transformation of entanglement.

4 Quantum cryptography

If one day, the impressive potential arising from the application of quantum mechanical principles to computer science will be achieved and building a practical quantum computer will no longer be a distant goal, perhaps the most spectacular impact will be seen on cryptography. The consequences on the security of cryptographic techniques and protocols can be both destructive and constructive. It is the aim of this section to show these effects, together with the current status of quantum cryptography.

4.1 Private-key systems

Historically, the first methods used to exchange secret information belong to the class of *private-key* systems. The two parties wishing to communicate have to meet and agree on a set of secret keys, subsequently used to encrypt and decrypt any messages between them. The *one-time pad* cryptosystem is an example from this category. Its security rests on several assumptions: each key is used only once, nobody is eavesdropping while the keys are randomly generated and agreed upon and finally, each of the two participants in the protocol is able to securely store its copy of the pad containing the secret keys. If these requirements are met, then the system is guaranteed to be totally foolproof. However, there is one major drawback all private-key systems share: the distribution of keys. While meeting face to face in order to distribute keys is not unconceivable for inter-governmental communications, it is certainly out of the question for commercial transactions over the Internet, for example.

4.2 Public-key systems

Public-key systems were invented exactly to address the problem of securely distributing cryptographic keys, without the communicating parties having to meet. In contrast to the symmetry of private-key systems, in which the same key is used for encryption and decryption, public-key systems use a pair of asymmetric keys. One is public, and therefore can be widely distributed, and

the other is private and must be kept secret. A message encrypted using the public key can only be decrypted using the associated private key. This is achieved using the mathematical concept of a *trapdoor function*. A trapdoor function is some mathematical procedure that is easy to compute, yet very hard to invert (for this reason they are sometimes called *one-way* functions), unless you have access to a special *key* that can unlock the *trapdoor*. Computing a trapdoor function corresponds to encrypting a message using the public key. To break such a code, one is forced to invert the trapdoor function without knowledge of the private key, while the intended recipient of the encoded message can use his or her private key in order to easily decipher the encrypted message.

4.2.1 Diffie-Hellman algorithm

The first practical public-key algorithm was devised by Diffie and Hellman in 1976 [69]. It was a continuation of Ralph Merkle's work on this track [138]. The Diffie-Hellman procedure does not completely depart from the symmetry of private-key systems. Alice and Bob, the prototypical participants in any cryptographic protocol, exchange their public keys in order to construct (based also on their secret keys) a common session key, which they can use to encode and decode messages between them, just like in a normal private-key system. From this point of view, one can look at the Diffie-Hellman procedure as a safe way of distributing private keys using only public communications channels. Reconstructing one of the secret keys, given knowledge of the corresponding public key, amounts to solving a *discrete logarithm* problem, which is very hard to calculate if the numbers involved are all sufficiently large. The security of the Diffie-Hellman algorithm depends on this fact.

4.2.2 RSA algorithm

But the most successful public-key system in use today is based on the RSA algorithm developed in 1977 [156]. Unlike the Diffie-Hellman algorithm, it is a genuine public-key system, in the sense that anyone can send an encrypted message to Alice using only her public key. Alice will then use her private (secret) key in order to decode the message. The RSA algorithm is also more versatile, offering the possibility to digitally sign a document due to the interchangeability between the public and private keys in this algorithm. Its security ultimately rests on the difficulty of factoring large numbers in a reasonable amount of time, a problem thought to be intractable, although **nobody was able to prove that it is not in P** . On the other hand, **nobody was able to prove it as an NP -complete problem either**.

The lack of a polynomial time solution to the factoring problem means that RSA encryption codes are safe for the time being, provided long enough keys are used. More precise, a 129-digit number used as the public key to encrypt a message by applying the RSA algorithm was already factored in 1994. However, this result was possible only after a gargantuan computational effort that lasted about eight months and involved some 1600 computers distributed over the Internet. Improvements in factoring technology made possible a much quicker factorization of a 130-digit RSA key two years later. Replacing the quadratic sieve with the number field sieve as the factorization method sped-up the computation approximately by a factor of 5. The number field sieve [126] is currently the best-known classical algorithm for factoring large integers. Its running time, while still super-polynomial, is sub-exponential in the size (number of digits) of the number being factored.

4.3 Cracking RSA codes

There are two useful remarks we can make here. First, the power of parallelism should not be underestimated. Lenstra, who was the first together with several colleagues to factor the ninth Fermat number $2^{2^9} + 1$ (155 decimal digits) and was involved in cracking both RSA-129 and RSA-130 thinks that RSA codes depending on 512-bit numbers (corresponding to 155 decimal digits) are within the reach of cypherpunks, if they could get hundreds of thousands of machines computing for them (see [42] page 167).

Secondly, since the security of RSA codes is based upon unproven assumptions, advancements in factoring algorithms are always possible. Coupled with technology improvements that allow ever faster computers to be built, this will force people to use longer and longer keys in order to keep the RSA systems secure. However, unless a polynomial time algorithm is discovered, factoring a 2000-digit number would take more than the entire life of the Universe, assuming that every particle in the Universe was a classical computer working at full speed (see [181] page 129). Even messages encrypted using a 400-digit key are well beyond the decryption capabilities of modern computers.

The cryptographic landscape, as it is perceived today, would change dramatically though if large-scale quantum computers could be built. This is because the algorithms discovered by Peter Shor in 1994 for solving the discrete logarithm problem and factoring numbers require only polynomial time on a quantum computer. Recall that the security of Diffie-Hellman procedure and RSA algorithm respectively, rely exactly on the presumptive intractability of these two problems on a classical machine. Given the current spread of public-key systems (especially RSA) and the increasing role they are expected to play as electronic commerce over the Internet will grow, the feasibility of a quantum computing device capable of executing Shor's algorithms would deliver a devastating blow to public-key cryptography. We will sketch the main features of Shor's factorization algorithm in the next section, in the more general context of quantum algorithms.

It appears that computing devices with quantum mechanical capabilities would render public-key systems useless, bringing the problem of distributing cryptographic keys through insecure communications channels back into attention. It turns out that the principles of quantum mechanics, which reactivated the problem in the first place, are also able to solve it. The answer comes from quantum cryptography, a discipline that has flourished initially independently of quantum computing, but now as an integral part of it, since both depend on the same technology and the same conceptual foundations.

4.4 Quantum key distribution

The birth of quantum cryptography may be assimilated with Stephen Wiesner's ideas in the late 1960s to use quantum mechanics in order to produce unforgeable money [180]. Although not very practical, his ideas stimulated other researchers and eventually Bennett and Brassard [18] were able to develop a protocol (generally referred to as BB84) for implementing a system of quantum key distribution. Their method employed quantum mechanical principles to exchange secret keys securely over public channels. So, unlike the public-key systems in use today, the unbreakability of quantum key distribution seems to be guaranteed by the very laws of physics (quantum mechanics, in this case).

The quantum key distribution scheme is amenable to any kind of physical realization, but photon polarizations offer a convenient way to explain and implement it. Suppose Alice and Bob wish to construct a key consisting of a random sequence of bits, known only to them, that they will subsequently use to encrypt and decrypt messages between them, using any private-key algorithm. Alice chooses a random set of bits that she encodes in either the rectilinear polarization (horizon-

Alice	0	1	1	0	0	1	0	0	0	1	0	0	1	1	0
	×	+	×	+	+	+	×	×	+	+	×	+	×	×	×
	↗	↑	↖	→	→	↑	↗	↗	→	↑	↗	→	↖	↖	↗
Bob	+	+	×	+	×	×	+	×	×	+	+	+	×	+	×
	1	1	1	0	0	1	1	0	1	1	0	0	1	1	0
key		1	1	0				0		1		0	1		0

Figure 5: Quantum key distribution in the absence of eavesdropping.

tal/vertical) or the diagonal polarization ($45^\circ/135^\circ$). The choice of polarization orientation must be random too (see Figure 5). Alice then sends the stream of photons (each carrying a bit of information) to Bob over an open quantum communication channel. Upon receipt of the photons, it is Bob's turn to choose an orientation for the device used to measure the direction of polarization of the incoming photons. According to the choices made between rectilinear and diagonal polarization measurement for each photon, Bob will extract a certain set of bits from the observed photons. In the last step of their protocol, Alice and Bob enter into a public communication. Alice divulges the polarizer orientations she used to encode the bits, while Bob informs Alice about the measurement basis he used to decode each photon received. Those bits for which both have chosen the same orientation will form their shared secret key. The others will simply be discarded. Note that the exact sequence of bits exchanged between Alice and Bob is irrelevant, as in the end they, and only they, come to learn the identity of a common subset of the bits, without having to reveal them to each other or to the outside world.

The above scenario is only possible if the encoding, decoding and transmission steps are error-free, and perhaps more important, there is no eavesdropping. If Eve, the prototypical eavesdropper, intercepts the photons on their way from Alice to Bob and wishes to gain some information about the secret key, she is forced to measure them, just as Bob would do. Unfortunately for her, the principle of quantum mechanics regarding measurements is not on her side and the quantum state of those photons for which Eve chooses an incorrect measurement basis will be inevitably disturbed. After discarding the irrelevant bits, Alice and Bob can randomly test some of the remaining bits to reveal the presence of any potential eavesdropper. For each bit tested, the probability of that test revealing Eve's presence is $1/4$. Thus, the probability of detecting eavesdropping can be made arbitrarily close to 1, by testing a sufficiently large number of bits. Since the tested bits can no longer be part of the key, there is a trade-off here between the desired level of security and the amount of quantum and classical communication required between Alice and Bob. Bennett and Brassard have addressed this issue by proposing alternate methods of testing, like parity checking, which leaves open the possibility for the set of bits involved in the test to still be included in the private key. Quantum key distribution is a good example of a case where quantum mechanical postulates seeming very restrictive and imposing severe limitations (like measurement and no-clonability, in this case) can actually be used in a constructive way and have important practical applications.

4.4.1 Experimental demonstrations

The first practical demonstration of their quantum key distribution protocol was performed by Bennett and Brassard in 1989 [19] over a distance of 32 cm. In spite of the insignificant distance, the experiment was very important as it showed that quantum cryptography is viable, but it also pointed out some difficulties. Trying to hide her presence behind the noise in the quantum channel and the imperfections of Bob's photon detectors, a clever eavesdropper could choose to measure

only some of the passing photons, according to the idea that partial information about the key is better than no information. To cope with such low levels of eavesdropping, Bennett, Brassard and Robert [22] have proposed the method of *privacy amplification*, a mathematical technique based on the principle of hashing functions that magnifies Eve's uncertainty over the final form of the key. There is also an efficient way of applying privacy amplification to the problem of reducing string oblivious transfer to bit oblivious transfer [39].

The ensuing rapid progress in quantum cryptography led to several experimental demonstrations of quantum key distribution over tens of kilometers of fiber-optic cable [84, 105, 106, 133, 142, 108]. Townsend and Marand at British Telecom [133] have even managed to convey their cryptographic signals alongside simulated telephone traffic carried on the same fiber-optic cables. Taking things one step closer to reality, Townsend also demonstrated how quantum keys could be securely distributed to many different receivers over a passive fiber-optic network. Notable is also the error-cancellation technique used by the team from University of Geneva [142].

However, using optical fibers as the transmission medium for photons has its limitations. Because the error rate increases with the length of the optical cable, quantum key distribution beyond 100 km is inconceivable today³. One proposed solution to the range limitation is to establish a chain of keys between Alice and Bob, but this may leave the signals vulnerable at the repeater stations. The distance limitation of cryptography could be lifted completely if genuine quantum repeater stations based on teleporting quantum states could eventually be built.

Finally, another exciting possibility currently under exploration is to perform quantum key distribution directly through the atmosphere using lasers and satellites, with no need for special optical fibers. The main problem in this case is to combat the disturbances caused by atmospheric turbulences and background radiation. The hope is that from a large number of raw bits exchanged, a reasonable number of error-free key bits could be distilled. Although free-space experiments have got further to go before they catch up with the fiber-optic demonstrations, progress along this path has been achieved too. From 30 cm initially demonstrated by Bennett and collaborators in 1991 [17], gradual improvements have made possible free-space quantum key distribution over 10 km during daylight and at night [109].

Thus, the remarkable experimental progress shown in quantum key distribution imposes quantum cryptography as the most promising practical application of quantum information theory, at least in the short term. Even at the current stage of development, quantum key distribution schemes using fiber-optic technology are sufficiently advanced to "wire the financial district of any major city" ([182] page 155). There is reason to believe that it will not take long before such technologies will become commercially viable.

4.5 Quantum solutions to discreet decision problems

So far we have seen how quantum key distribution in conjunction with a guaranteed secure classical cryptosystem (like the one-time pad, for instance) make practical classical communications as secure as our current knowledge of physics allows. An important question is whether quantum cryptography can have other applications beside quantum key distribution. Classical cryptography is able to offer solutions to a wide variety of situations classified as *discreet decision* problems. In these situations, discretion is vital to achieving agreements. Examples include negotiating arms treaties, forming business partnerships and organizing mergers. These applications are amenable to classical cryptographic solutions involving public-key systems, but as we have seen, they are based on unproven assumptions about the difficulty of factoring large numbers and other related

³Note that conventional repeaters cannot be used due to the no-cloning theorem.

problems. What quantum cryptographers were looking for, was a totally secure system, guaranteed by the laws of physics.

Research in classical cryptography has shown that solutions to such discreet decision problems can be constructed from basic cryptographic building blocks, like the notion of *oblivious transfer* [154] or 1-out-of-2 oblivious transfer, an idea foreseen in a quantum guise by Stephen Wiesner [180]. Later, Claude Crépeau and Joe Kilian have demonstrated that oblivious transfer can be used as a building block for solving two-party problems requiring discretion [58]. In turn, to provide totally secure quantum oblivious transfer, one would need a secure form of bit commitment. Consequently, much of the research effort in quantum cryptography in the early 1990's was devoted to finding a protocol for quantum bit commitment that is absolutely and provably secure. That result (known as BCJL after the authors' names) was reported in 1993 [38] and became the foundation for numerous applications in quantum cryptography, pertaining to discreet decision making.

The surprise came in 1995 when Dominic Mayers discovered how Alice could cheat in the BCJL bit commitment protocol by using EPR particles [136]. Furthermore, he proved [137] that it would be possible for Alice to cheat in *any* protocol for quantum bit commitment. An intuitive explanation is that the information she sends Bob describing the *safe* for her bit must give nothing away about the *committed bit*. Consequently, regardless of the particular bit commitment scheme employed, the quantum states of the safe containing either 0 or 1 must be very similar (if not identical) since otherwise Bob would be able to discern the difference and gain knowledge about the committed bit prematurely. But the very fact that the two states are virtually the same gives Alice the possibility to keep her options open and postpone her *commitment* for later on. Although in their 1996 review paper of quantum cryptography, Brassard and Crépeau [37] argued that for the time being the practical implications of the flaw discovered in the quantum bit commitment protocol are minimal, the weakness definitely affected the entire edifice of quantum cryptography built upon quantum bit commitment.

4.6 Entanglement-based cryptography

However, analogue to the two-fold influence of quantum mechanics upon cryptography, entanglement can also be used in a beneficial way, enhancing methods of key distribution. Inspired by EPR experiments designed to test Bell's inequality, Artur Ekert thought of a way of using entangled pairs for distributing cryptographic keys by quantum means [75]. In his scheme, Alice and Bob receive entangled particles from a central source and perform independent measurements upon them. For each measurement, the orientation is chosen at random among three possibilities. The presence of any potential eavesdropper can be revealed when Alice and Bob publicly confront the results they got for the measurements in which they adopted different orientations. If the original EPR pairs were untampered with, then the strength of their correlations must exceed anything that is possible classically. These correlations will be disrupted if someone attempts to make measurements on the particles, before they arrive at the legitimate receiver.

So the security of Ekert's quantum key distribution scheme rests exactly on the dismissal of the hidden-variable theory. Eve's only hope is that entangled particles might carry hidden information about how they will behave in subsequent measurements. But she cannot elicit any information from the transiting particles simply because there is no information encoded there. The information about the secret key has yet to come into being, once Alice and Bob perform their measurements.

Ekert's entanglement-based scheme also offers a couple of potential advantages over the original single photon protocol invented by Bennett and Brassard. The first refers to the possibility of storing cryptographic keys securely, while the second involves the issue of privacy amplification.

The limitations of the classical privacy amplification based on hashing algorithms are overcome in the quantum privacy amplification technique, developed in 1996 by a group of researchers including David Deutsch, Artur Ekert and Richard Jozsa [66]. The quantum procedure, which is applicable only to entanglement-based quantum cryptography, can be repeatedly applied to impurely entangled particles to cleanse them of any signs of tampering by Eve. The *entanglement purification* process is actually an extension of previously published work [21] and needs only some simple quantum logic.

However, these advantages of entanglement-based cryptography are rather theoretical at the moment because storing entangled particles is only possible for a fraction of a second as yet, and entanglement purification depends on quantum computational hardware that, although simple, has yet to be built. This, of course, assuming that entanglement is indeed a real physical resource that can be harnessed for our computation or communication purposes. Although Ekert and Rarity [77] devised a plan for implementing a practical method of entanglement-based cryptography and Nicholas Gisin's team in Geneva reported an experimental demonstration of quantum-correlations over more than 10 kilometers [175], these results are nowhere near the remarkable progress achieved by quantum key distributions using the original BB84 protocol, which is well within the capabilities of current technology.

Work has also been carried out to develop efficient cryptographic protocols based on noisy channels [57] and, more recently, to comparatively assess the potentials of classical and quantum protocols for key agreement in various situations [89, 90, 88].

5 Quantum algorithms

Quantum computing is a fundamentally novel paradigm with respect to both the "hardware" and the "software" employed. The hardware in a quantum computer concerns the physical realization of the qubits organized into a memory register and the quantum logic gates acting on them. Designing a quantum algorithm to solve a certain problem boils down to the choice of the quantum gates, that is unitary transformations, which, when chained together in a proper way, evolve the initial state of the quantum memory register into the desired final state. Thus, we can say that the LOAD-RUN-READ operational cycle characteristic to a classical computer is replaced by a PREPARE-EVOLVE-MEASURE cycle in a quantum computer.

Quantum parallelism The "programming techniques" employed to achieve the desired evolution are essentially different from their classical counterparts. This is due to the necessity of playing by the rules imposed by quantum mechanics, which severely restrict the applicability of classical algorithmic design methods. The main feature of a quantum computer, used in any quantum algorithm, is *quantum parallelism*. Quantum parallelism refers to the capability of a quantum computer to evaluate a function $f(x)$ for exponentially many different values of x in the time it takes a classical computer to evaluate the function for just one value. This is possible by loading the memory register with a superposition of all possible input values x and then apply the appropriate unitary transformation that will evolve this state into a superposition of all function values $f(x)$. The enormous potential of a quantum computer to outperform a classical machine lies in the massive parallelism it offers "within a single piece of hardware", as Berthiaume and Brassard put it [28].

Extracting the answer However, a direct observation of the quantum memory register will not yield more information than is possible to obtain using a classical computer. Any measurement

attempt will collapse the superposition and reveal only one of the answers, without even knowing which one beforehand. Therefore, quantum parallelism alone does not justify all the trouble of going into quantum computing. Additional techniques seem necessary in order to exploit quantum parallelism and make a quantum computer useful. Fortunately, quantum mechanics has also the resources to achieve this task: interference and entanglement. The heart of any quantum algorithm is the way in which it manipulates entanglement and interference between various computational paths, so that the final measurement will yield desired results with high probability. Generally, the design of a quantum algorithm is focused on how interference can constructively recombine different alternatives in a superposition to strengthen the amplitude of solutions, while non-solutions will interfere destructively, canceling each other. According to the specific technique employed to achieve this sort of manipulation as well as their area of applicability, we can identify three main classes of quantum algorithms which provide an advantage over known classical algorithms: quantum algorithms based upon some kind of Fourier transform, quantum search algorithms and quantum simulations.

5.1 Algorithms using quantum Fourier transforms

The discrete Fourier transform can be brought about in a quantum mechanical context by defining a linear transformation on n qubits whose action on the computational basis states $|j\rangle$, $0 \leq j \leq 2^n - 1$, is described below:

$$|j\rangle \xrightarrow{QFT} \frac{1}{\sqrt{2^n}} \sum_{k=0}^{2^n-1} e^{\frac{2\pi i j k}{2^n}} |k\rangle. \quad (8)$$

Expressing the effect of applying this Quantum Fourier Transform on a superposition of the computational basis states, in the form

$$\sum_{j=0}^{2^n-1} x_j |j\rangle \xrightarrow{QFT} \sum_{k=0}^{2^n-1} y_k |k\rangle, \quad (9)$$

corresponds to a vector notation for the usual discrete Fourier transform, in which the complex numbers x_j are transformed into the complex numbers y_k .

5.1.1 Deutsch's algorithm

Under various forms, quantum Fourier transforms play an essential role in many algorithms, bringing about the necessary interference in order to find common properties of all the values in a superposition without having to reveal any of the individual values explicitly. The first algorithm to exploit this idea was devised by David Deutsch as an example of how quantum parallelism backed by interference can "beat" a classical computer [61]. Given a function $f : \{0, 1\} \rightarrow \{0, 1\}$ Deutsch presented a quantum algorithm able to compute $f(0) \oplus f(1)$ in a single evaluation of the function f . Beside computing $f(0)$ and $f(1)$ in quantum parallel using a superposition of the two inputs, Deutsch was able to encode the value of the function $f(x)$ in a phase factor $(-1)^{f(x)}$. This, in turn, enabled a quantum mechanical interference between the phase factors to reveal the desired joint property of $f(0)$ and $f(1)$. It is important to note that Deutsch's algorithm can only answer the question whether $f(0)$ equals $f(1)$ or not, without giving any information about either function value individually.




Figure 6: Quantum circuit implementing the Deutsch-Jozsa algorithm.

5.1.2 The Deutsch-Jozsa algorithm

Deutsch and Jozsa [67] generalized this problem to the n -bit case, allowing the domain of f to be $[0, 2^n - 1]$ and determining whether f is constant for all values of x or perfectly balanced between 0 and 1 in just one evaluation of the function f . The quantum circuit implementing the general Deutsch-Jozsa algorithm is depicted in Figure 6. The Walsh-Hadamard transform $H^{\otimes n}$ is a generalization of the single-qubit Hadamard gate H and corresponds to n Hadamard gates acting in parallel on n qubits. When applied to n qubits, all prepared in the $|0\rangle$ state, the Walsh-Hadamard gate creates an equally weighted superposition of all integers in the range $[0, 2^n - 1]$. This makes it widely used in the initial step of virtually any quantum algorithm. The application of the Hadamard gate on the single qubit $|1\rangle$ gives the superposition $1/\sqrt{2}(|0\rangle - |1\rangle)$, responsible for encoding the results of the function evaluation in the amplitude of the corresponding term in the superposition state $|\Psi\rangle$:

$$|\Psi\rangle = \sum_{x_1 \cdots x_n} \frac{(-1)^{f(x_1 \cdots x_n)} |x_1 \cdots x_n\rangle}{\sqrt{2^n}} \left[\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right]. \quad (10)$$

In the final step of the computation, another Walsh-Hadamard gate acts on the first n qubits interfering the terms in the superposition:

$$|\Psi'\rangle = \sum_{z_1 \cdots z_n} \sum_{x_1 \cdots x_n} \frac{(-1)^{x_1 z_1 + \cdots + x_n z_n + f(x_1 \cdots x_n)} |z_1 \cdots z_n\rangle}{2^n} \left[\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right]. \quad (11)$$

A closer look to the amplitude for the state $|0\rangle^{\otimes n}$, namely

$$\sum_{x \in \{0,1\}^n} \frac{(-1)^{f(x)}}{2^n}$$

shows that it is possible to discern whether f is constant or balanced by measuring the first n qubits. If f is constant, then the above amplitude is $+1$ or -1 depending on the constant value $f(x)$ takes. However, in either case, all the other amplitudes must be zero because $|\Psi'\rangle$ is of unit length, by definition. Therefore, if the measured qubits are all 0, we know that f must be constant. On the other hand, when f is balanced, the positive and negative contributions to the amplitude for $|0\rangle^{\otimes n}$ cancel each other, giving an amplitude of zero. This means that at least one of the measured qubits must yield a result of 1.

The original algorithms for Deutsch's problem and the Deutsch-Jozsa problem have been substantially improved subsequently by Cleve, Ekert, Macchiavello and Mosca [54]. The original

algorithm of Deutsch, for example, only worked probabilistically, giving a meaningful answer only half the time. The presentation above is based on the improved versions of these algorithms.

5.1.3 Instances of quantum Fourier transforms

Other researchers tried to extend the work of Deutsch and Jozsa by developing variants of their balanced versus constant problem. Berthiaume and Brassard, for instance, studied the original Deutsch-Jozsa problem in an oracle setting [28]. Dan Simon developed a variant that showed how a quantum computer could efficiently decide whether a function is 1-to-1 or 2-to-1 [165]. Although only periodic 2-to-1 functions were allowed, the interesting thing about Simon's algorithm was that for such functions it was able to pluck out their periodicity by exploiting a simple kind of Fourier transform. But the first paper to explicitly show how Fourier transforms can be implemented on a quantum computer belonged to Bernstein and Vazirani. In their paper about quantum complexity theory [26] they also showed how to sample from the Fourier spectrum of a Boolean function on n bits in polynomial time on a Quantum Turing Machine. Nevertheless, as it was later realized, the essence of the Fourier transform idea was actually already hidden within the Deutsch-Jozsa quantum algorithm solving the constant versus balanced problem. Indeed, the Walsh-Hadamard transformation applied in the last step of the algorithm can be intuitively assimilated with a quantum Fourier transformation which generates an interference between all computational paths.

5.1.4 Shor's factorization algorithm

All these developments progressively enlarged the scope of what was possible using quantum algorithms, paving the way for the most spectacular result as yet, namely Peter Shor's algorithm for factoring large integers and computing discrete logarithms efficiently on a quantum computer [164]. Of great inspiration to Shor was the work of Dan Simon [165] and Bernstein and Vazirani [26]. As a number theorist, Peter Shor was well aware of the relation between factoring a number n and calculating the orders (or periods) of numbers *modulo* n . He was hoping to build upon Simon's paper about finding periodicities in 2-to-1 functions and devise an efficient method to compute the orders for functions of the form $f_{x,n}(a) = x^a \bmod n$. In classical complexity theory it has been long known that finding such orders when the modulus n gets very large is as hard as factoring n . In the same paper that offered a polynomial-time algorithm for primality testing, Gary Miller showed that the problem of finding orders is computationally equivalent to the problem of factoring integers [140]. But Shor was counting on the quantum Fourier transform to find orders efficiently using quantum mechanics. Once an even number period r is found for $x^a \bmod n$ by varying the value of x , $x^{r/2} - 1$ and $x^{r/2} + 1$, respectively, have a good chance of sharing a common divisor with n . A sketch of Shor's quantum algorithm for factoring integers is given below.

1. Set up a quantum memory register partitioned into Register 1 and Register 2. Load Register 1 with a superposition of all possible integers by means of a Walsh-Hadamard gate. Load Register 2 with zeros.
2. Pick a random integer x which is coprime with n and evaluate the function $x^a \bmod n$, in quantum parallel, for all terms in the superposition, such that each number a in Register 1 is entangled with the corresponding result, placed in Register 2.
3. Observe the state of Register 2 to be some integer k . Because of the entanglement between the two registers, this measurement will also have the effect of projecting out the state of Register 1 to be a superposition of just those values of a such that $x^a \bmod n = k$.

4. Compute the Fourier transform of the post-measurement state of Register 1.
5. Measure the state of Register 1 to sample from the Fourier transform and obtain some multiple of $2^q/r$, where q is the number of qubits in Register 1. Use a continued fraction technique to determine the period r .
6. Obtain the factors of n by computing $\gcd(x^{r/2} - 1, n)$ and $\gcd(x^{r/2} + 1, n)$.

Analysis Note that Shor's algorithm is probabilistic. Several things can go wrong: the period of $f(a) = x^a \bmod n$ is odd, the value sampled from the Fourier transform is not close enough to a multiple of $2^q/r$, the continued-fraction expansion yields a factor of r and not the period itself, $x^{r/2} - 1$ or $x^{r/2} + 1$ is a multiple of n . Nevertheless, Shor showed that few repetitions of this algorithm reveal a factor of n with high probability, thus providing a bounded probability polynomial time algorithm for factoring numbers on a quantum computer. A condition for keeping the complexity of the algorithm at a polynomial level was to devise an efficient quantum implementation for the Fourier transform in step 4. In his paper, Shor describes a way of constructing the quantum Fourier transform with base 2^m using only $m(m+1)/2$ gates. Subsequent work by Adriano Barenco and colleagues at Oxford [9] showed that the circuit required for computing the QFT can be simplified by making certain approximations, thus even becoming computationally less demanding than the procedure for calculating the powers stored in the second register in step 2.

There is also an important observation that has to be made about step 3 in the algorithm described above. The possibility of a measurement made on the contents of the second register before carrying out the Fourier transform was actually hinted at in a later paper [8]. In Shor's original paper, he suggested making the measurement only after the Fourier transform on the first register. This may look quite intriguing from a classical computational point of view, since the values of the function $f(a) = x^a \bmod n$ computed in Register 2 seem not to be referred to again. This apparent redundancy illustrates very well the differences between conventional and quantum mechanical programming techniques and highlights some of the subtleties of working with entanglement in superpositions. The entanglement between Register 1 and Register 2 ensures that only those amplitudes corresponding to numbers a having the same image $f(a)$ will be able to interfere when the Fourier transform on Register 1 is invoked. The result will be a superposition of Fourier transforms of a set of functions having the same period r , but which do not interfere with each other. Therefore, the measurement in step 3 can be skipped entirely. More generally, Bernstein and Vazirani [26] showed that measurements in the middle of an algorithm can always be avoided.

Thus, Shor's quantum algorithm for factoring integers relies on quantum parallelism to create a superposition of values of the periodic function $f_{x,n}(a)$, relies on entanglement and the Quantum Fourier Transform to create the desired interference effect between solutions (integer multiples of $1/r$) and non-solutions (numbers that are not integer multiples of $1/r$) and finally relies on measurement to project out a multiple of the inverse of the sought-after period r . The rest are techniques from classical number theory. Another useful application of the quantum fast Fourier transform was found by Abrams and Lloyd [1]. They managed to devise a new polynomial time quantum algorithm that finds eigenvalues and eigenvectors of certain matrices for which all known classical algorithms require exponential time.

Generalization The usefulness of quantum Fourier transforms has prompted the development of a generalized theory of quantum Fourier transforms involving some technical ideas from the theory of finite groups. This line of thought culminated in Kitaev's discovery of a method to solve the

Abelian stabilizer problem [118] and the generalization to the hidden subgroup problem [98]. The Deutsch-Jozsa algorithm, Shor's algorithms and related exponentially fast quantum algorithms can all be viewed as special cases of this algorithm. **It would be interesting to know whether other problems of practical importance can be accommodated to fit this general framework nicely.**

Recently, steps have been taken to develop new techniques that are not based on quantum Fourier transforms, but still provide an exponential algorithmic speed-up relative to a classical computer. Childs et al. [47] show how the quantum walk model proposed in [79] can be exploited to construct an oracular (*i.e.* black box) problem in which a graph whose structure is provided in the form of an oracle can be traversed exponentially faster than is possible by any classical algorithm.

5.2 Quantum search algorithms

Another methodology to construct useful quantum algorithms is based on the idea to transform a superposition quantum state in such a way as to amplify the values of interest at the expense of the amplitudes of the other terms in the superposition. This will give the solutions a higher probability of being measured in the end. The most representative exponent of this class is Grover's unstructured search algorithm [94].

5.2.1 Grover's search algorithm

Formally, the unstructured search problem can be defined as finding some x in a set of possible solutions such that a certain statement $P(x)$ is true. In addition, no assumption is to be used about the structure of the search space and the statement P . As expected, Grover's algorithm relies on the same quantum mechanical principles that give quantum algorithms in general the upper hand over corresponding classical algorithms. We refer to quantum parallelism, entanglement and interference, as can be seen from the description of Grover's algorithm given below.

1. Prepare two quantum registers, the first containing a superposition of all possible input values $x_i \in [0..2^n - 1]$ and the second one set to 0.
2. Compute $P(x_i)$ in quantum parallel for all the values in the first register, creating a superposition of results in the second one, which is now entangled with the first.
3. Repeat approximately $\frac{\pi}{4}\sqrt{2^n}$ times
 - 3.1 Change the sign of the amplitude for the state x_j such that $P(x_j) = 1$.
 - 3.2 Invert all the amplitudes about the average. This will increase the amplitude of the target state, while all the other amplitudes will be diminished imperceptibly.
4. Read the result. If Register 2 is 1 (the value representing *True*), Register 1 will contain the sought-after value x .

Analysis Clearly, the most important step in Grover's algorithm is the amplitude amplification operation performed in step 3. At the first glance, it might look surprising how we can change the sign only for the target state without knowing it beforehand. However, if we evaluate the predicate P by means of a gate array U_P performing the transformation $|x, b\rangle \rightarrow |x, b \oplus P(x)\rangle$, we can apply U_P to the superposition

$$\frac{1}{\sqrt{2^n}} \sum_{x=0}^{2^n-1} |x\rangle$$

and choose $b = 1/\sqrt{2}(|0\rangle - |1\rangle)$. This way, we will end up in a state where the sign of all x with $P(x) = 1$ has been changed and b is unchanged. Grover also showed that the inversion about the average can be accomplished efficiently by decomposing it into $O(n)$ elementary quantum gates with the aid of the Walsh-Hadamard transform.

Grover's algorithm is another good opportunity to point out the fundamental differences between classical and quantum programming techniques. Many classical algorithms rely on indefinite repetitions of the same procedure to keep improving the results. In contrast, repeating a quantum procedure may improve results only up to a certain point, after which the results will get worse again. The reason is that quantum procedures are unitary transformations which can be interpreted as rotations of the vector state in a complex space. Thus, the repeated application of a quantum transform may rotate the initial state closer and closer to the target state, eventually getting past it farther and farther away. Therefore, the optimal number of iterations of a quantum procedure is an important design issue. Boyer et al. [36], who provide a detailed analysis of Grover's algorithm, show that for a single solution x_0 such that $P(x_0)$ is true, $\frac{\pi}{8}\sqrt{2^n}$ iterations of step 3 above will bring the failure rate to 0.5. After another $\frac{\pi}{8}\sqrt{2^n}$ iterations, the failure rate will drop to $1/2^n$. If we keep iterating for another $\frac{\pi}{4}\sqrt{2^n}$ times though, the final measurement is almost guaranteed to give us a non-solution.

The number of repetitions of step 3 also determines the complexity of Grover's algorithm. It follows that on a quantum computer, the unstructured search problem can be solved with bounded probability of error within $O(\sqrt{N})$ evaluations of P , where N is the size of the search space. This represents only a quadratic speed-up relative to the best possible classical algorithm, which is far less impressive than the exponential speed-up achieved by Shor's factoring algorithm, for example. Still, the importance of quantum search algorithms is justified by the fact that searching heuristics have a wider range of applications than problems solved using the quantum Fourier transform and if adaptations of the quantum search algorithm are also taken into account, the range of problems that can benefit from them is even broader.

Extensions Biron et al. [31] showed how Grover's search technique can be used with arbitrary initial amplitude distributions, while still maintaining the overall $O(\sqrt{N})$ complexity. This means that Grover's algorithm can be used as a subroutine in other quantum computations. Grover's algorithm has also been combined with Shor's algorithm in order to perform quantum counting, that is to determine the number of solutions and the optimal number of iterations [40]. Thus, if the search space contains S solutions, then a quantum computer takes only $O(\sqrt{N/S})$ steps to find one of them. Grover himself extended his algorithm in two ways: to show that certain search problems that classically run in $O(\log N)$ can be solved in constant time on a quantum computer, and to achieve quadratic speed-up for other non-search problems such as computing the mean and median of a function [95]. The intrinsic robustness of Grover's quantum search algorithm in the presence of noise is evaluated in [149].

5.2.2 Heuristic quantum search algorithms

For completely unstructured searches, Grover's algorithm is optimal [16, 36, 186]. However, there are also search problems, like constraint satisfaction problems such as β -SAT or graph colorability,




Figure 7: Lattice of variable assignments.

in which information about the search space and statement P can be exploited for heuristic algorithms that yield efficient solutions for some problem instances. The techniques of Grover's search algorithm can be used to construct quantum analogs with quadratic speed-up to classical heuristic searches [46, 40]. But there is also hope that for certain structured problems a speed-up greater than quadratic is possible by using specific properties of quantum computation to exploit problem structure and not just merely construct quantum implementations of the classical algorithms. The work of Tad Hogg is representative for this approach [99, 100, 101, 102]. He has developed heuristic quantum search algorithms that exploit the inherent structure of constraint satisfaction problems in a distinctly non-classical way, which uses unique properties of quantum computation, like interference.

Constraint satisfaction problems Solutions to a generic constraint satisfaction problem lie in the space of assignments of m different values to n variables such that a set of l constraints are satisfied. Such a search space harbors a natural structure taking the form of a lattice given by set containment. Figure 7 shows the lattice structure formed in the assignment space spanned by two variables taking two possible values: 0 and 1. The sets in this lattice can be put in one-to-one correspondence with the standard basis vectors for a four-qubit quantum state space, such that a 1 in the binary sequence representing a basis vector corresponds to inclusion, while a 0 corresponds to exclusion of the respective element (variable assignment). The result will be a lattice of variable




Figure 8: Lattice of variable assignments in ket form.

assignments in ket form, depicted in Figure 8.

Moving the amplitude up Hogg's idea is to start with a single initial amplitude concentrated in the $|0 \cdots 0\rangle$ state and then iteratively move amplitude up the lattice, from sets to supersets and away from "bad sets" (*i.e.* sets that violate the constraints). Note the interesting particularity of this algorithm, to start differently from all the quantum algorithms presented so far (Deutsch-Jozsa, Shor's, Grover's) which all begin by computing a function on a superposition of all the input values at once.

Hogg devised two methods for moving the amplitude up the lattice [99, 100] by constructing a unitary matrix which maximizes the movement of amplitude from a set to its supersets. For moving the amplitude away from bad sets, Hogg suggests a general technique based on adjusting the phases of the bad sets and using interference to decrease the amplitude of sets that have bad subsets (due to cancellations), while the amplitude of good subsets will add to increase the amplitude of the encompassing superset. Depending on the particular problem to be solved, the choice of different policies for manipulating the phase changes will result in different cancellations obtained. The technique exploits the property that if a state violates a constraint, then so do all states above it in the lattice.

Efficiency The problem with Hogg's algorithms is one shared by heuristic algorithms in general. The use of problem structure is complicated enough not to allow for an accurate estimation of the probability of obtaining a solution from a single execution of the algorithm. Therefore, it is difficult to analyze the efficiency of Hogg's quantum search heuristics. The efficiency of classical heuristic algorithms is estimated by empirically testing the algorithm. Since a practical quantum computer on which Hogg's algorithms could be tested was not yet built (nor will it be in the near future), all we can do is simulate his quantum algorithms on a classical computer. Unfortunately, this incurs an exponential slowdown, thus making it feasible only on small cases. From the few

small experiments that have been done, the guess is that Hogg's algorithms are more efficient than Grover's algorithm applied to structured search problems, though the speed-up is likely to be only polynomial. Even so, for the computationally difficult problems to which they can be applied, a small polynomial speed-up on average is of significant practical interest.

5.3 Quantum simulations

If practical quantum computers will ever become a reality, then a class of tasks at which they could naturally outperform any classical machine is simulating quantum mechanical systems occurring in Nature. As the size (number of constituents) of a quantum system increases, the number of variables required to describe the state of the system grows exponentially. So, in order to store the quantum state of a system with n distinct components, a classical computer would need some c^n bits of memory, with the constant c depending upon the system being simulated and the desired accuracy of the simulation. Furthermore, calculating its evolution over time would require the manipulation of a huge matrix, involving $c^n \times c^n$ bits. On the other hand, a machine that worked by quantum means would intrinsically make a much more efficient simulator, requiring only a linear number of qubits. This is also the reason for which efficient simulations of a quantum computer on a classical machine are not known to exist.

Feynman was the first to hint at the idea that rich and complex dynamical evolutions of some quantum systems could be simulated resorting only to simple local interactions [80]. Feynman's idea was refined by Lloyd, who showed in 1996 that the logical operations available on a quantum computer could be marshaled to simulate the behavior of virtually any quantum system whose dynamics is determined by local interactions [129]. However, the problem of obtaining the desired information about the simulated quantum system from the information hidden in the wavefunction characterizing its state still remains. Thus, a crucial step in making quantum simulations useful is the development of systematic means by which desired answers can be efficiently extracted. As we have already seen, direct observation (measurement) is of little help due to the irretrievable loss of information incurred by the collapse of the wavefunction.

The impact of obtaining faster and more accurate simulations of systems in which quantum phenomena are important may materialize in significant advances in those fields. We enumerate here some of the areas that would benefit from such efficient simulations: molecular chemistry with direct applications in pharmaceutical industry, studying the behavior of liquids and gases, gaining insights about the nature of forces governing the behavior of nuclear particles, verifying the strength of various theories trying to explain superconductivity, especially at "high" temperatures. As Lloyd has noted, it might even be possible for quantum simulations to take advantage of a greater range of quantum hardware. Thus, for instance, decoherence, instead of being a liability and spend effort trying to neutralize its effects, could be turned to our advantage by using it to replicate the interaction of the simulated system with its environment. Finally, quantum simulation algorithms could also be employed as a general method to obtain insight into other quantum algorithms.

Shape of quantum circuits We are going to conclude this review of quantum algorithms with a general observation about the shape of the quantum circuits implementing them. As it became apparent with the unifying work of Richard Cleve, Artur Ekert and colleagues [54], quantum algorithms tend to have a rather similar structure. The quantum circuits used to describe their operation usually have a specific (problem dependent) unitary transformation in the middle, sandwiched between Hadamard or Fourier transforms. **It would be interesting to know whether more or less all quantum circuits will take this form.** A positive answer to this question




Figure 9: Quantum Toffoli (controlled-controlled-NOT) gate.

may induce the feeling that this structure offers only a rather limited range of opportunities.

6 Quantum complexity

Similar to classical complexity theory, quantum complexity theory is concerned with classifying the difficulty of various computational problems, grouping them into complexity classes according to the spatial (memory) and time resources needed by a quantum computer to solve those problems. Naturally, in this endeavor, of great interest is to compare the relative powers of the classical model of computation (represented by the Universal Turing Machine, for example) and the computational model based on quantum mechanics.

In the previous section, several examples of quantum algorithms were presented which perform better than any known classical algorithm confronted with the same computational problem. **But are quantum computers strictly more powerful than classical ones? Are there computational tasks which can be efficiently solved quantum-mechanically and yet no Turing machine is able to tackle them efficiently?** The answers to these questions is not known, despite the suspicions fostered by examples such as factoring, which suggest that the above questions can be answered in the affirmative. As a matter of fact, the initial excitement caused by an eventual positive answer has been tempered lately by some complexity theorists who now ask: **"How long until we can prove in the general case that a tight upper bound for quantum performance speedups is quadratic?"** [83]. This section is intended to offer a map containing the most important milestones in the development of quantum complexity theory.

6.1 Reversibility

The first step in evaluating the computational power of a quantum computer was to determine whether we can simulate a classical logic circuit using a quantum circuit. Bennett showed that any classical circuit can be converted into an equivalent reversible circuit, and moreover, that this conversion can be done efficiently [15]. One way to achieve this is to make use of a reversible gate known as the Toffoli gate, which is universal for classical computation and can be used as a basic building block in constructing reversible circuit models of computation [85]. Since the Toffoli gate is reversible (its inverse is itself), it can also be implemented as a quantum logic gate, depicted in Figure 9. Following its operation rules, the Toffoli gate can be seen as a controlled-controlled-NOT gate. Thus, we can immediately conclude that a quantum computer is at least as powerful as a

classical deterministic computer. Furthermore, quantum computers are also capable of performing any computation which a classical probabilistic computer may do by using a Hadamard gate to produce random fair coin tosses.

6.2 Quantum Turing Machines

In order to be able to prove theorems about further capabilities of quantum computers, researchers in quantum complexity theory tried to capture the essential features and principles underlying quantum computation into an abstract model, free of any particular implementation details. Since the Turing machine has proved to be the most influential among all models of classical computation, much of the research effort was focused on defining a quantum analogue of the Turing machine.

Charles Bennett had already shown in 1973 that a reversible Turing machine was a theoretical possibility [15]. Then, Paul Benioff realized the similarity between the operation of a reversible Turing machine and the dynamical evolution of an isolated quantum system, governed by the unitary evolution of the Hamiltonian in Schrödinger's equation. Thus, he devised a hypothetical quantum machine (system) whose evolution over time mimicked the actions of a classical reversible Turing machine [13]. However, Benioff's quantum mechanical implementation of a Turing machine did not incorporate any of the features responsible for the computational power attributed to quantum computers, such as superposition, interference and entanglement. For this reason, Benioff's design can do no more than a classical reversible Turing machine.

Although Feynman began by trying to simulate quantum systems on classical machines [80] (which is the opposite of what Benioff did), he also investigated the feasibility of a quantum mechanical computer based on a model of reversible computation developed by Toffoli and Fredkin [85]. His conclusion was that "the laws of physics present no barrier to reducing the size of computers until bits are the size of atoms, and quantum behavior holds dominant sway" [81]. Despite the fact that Feynman's model was an advance over Benioff's machine, it appears that its capabilities are also limited to mimic the operations of a general purpose computer, with no references to simulating quantum mechanics being made.

The first true quantum Turing machine (QTM) capable of exploiting genuine quantum mechanical effects was devised by David Deutsch in 1985 [61]. It called for a quantum mechanical processor responsible for controlling the read, write and shift operations performed by the head through quantum mechanical interactions. For each internal state in which the head can exist, the action of the QTM is specified by a set of quantum control rules taking the form of a Hamiltonian operator that specify the probability amplitude for the occurrence of all allowed state-to-state transitions. Deutsch's quantum computer also makes use of an infinitely long tape, on which both the program and the input data would be stored. Defining such an abstract quantum Turing machine was a key step in making it possible to study the computational power of quantum computers. It gave researchers in the field an important tool to address fundamental issues like universality, computability, provability and complexity.

6.3 Universality

The question of *universality* concerns whether a given machine can simulate all others and how efficiently can it do this. David Deutsch showed that it was possible to devise a Universal Quantum Turing Machine, but the simulation overhead was exponential in the running time of the simulated Turing Machine in the worst case. The efficiency of Deutsch's universal simulator for quantum Turing machines was improved by Bernstein and Vazirani [26] who were able to prove the existence of a Universal Quantum Turing Machine whose simulation overhead is polynomially bounded.

But quantum information theory and universality are also connected in a subtle and therefore unexpected way. The example presented in [143] shows that only a parallel approach can succeed in distinguishing between entangled quantum states. This result implies that no machine exists capable of simulating all possible computations. The notion of a Universal Computer is consequently a myth.

6.4 Computability

In terms of *computability*, we have already seen that anything computable by a classical computer is also computable by a quantum computer. But is the inverse of this statement also true? It is not difficult to envisage a classical Turing machine that simulates an arbitrary quantum circuit, if one does not care about efficiency. The simulation in [14] requires space, and therefore time, exponential in the number of qubits in the quantum circuit. Bernstein and Vazirani [26] have given a simulation that takes polynomial space, but exponential time. The lack of an efficient classical simulation of a quantum computer induced the idea that a quantum computing machine may be inherently faster. However, **no one has been able to prove that such an efficient simulation does not exist.**

So when it comes to the continuous evolution, described by Schrödinger's equation, of a quantum system, it seems that a quantum computer is no more powerful (with respect to computability) than a classical Turing machine. Nonetheless, a measurement operation causes the system to undergo a sudden, discontinuous transformation in which a superposed state will collapse onto one of the basis states. The outcome of the measurement depends on the probabilities given by the amplitudes of the components in the wave function. This intrinsic non-determinism in the behavior of any observed quantum system can only be faked by a classical machine through the use of a pseudo-random number generator. This observation led Deutsch to the conclusion that the notion of computability depends not on mathematical ideas about the nature of algorithms, but on the computational capabilities of physical systems. His point of view is reinforced by Rolf Landauer who stresses that information is a physical quantity: "Information is inevitably represented by real physical entities and is therefore tied to the laws of physics." (in [42] page 116). Consequently, Deutsch proposed the reformulation of the Church-Turing thesis in physical terms. In his book "The Fabric of Reality" he gives the following articulation of what he called *the Turing principle*: "There exists an abstract universal computer whose repertoire includes any computation that any physical possible object can perform." ([63] page 132).

There are also some theoretical results that seem to hint at the idea that quantum devices may be strictly more powerful than classical machines [44, 45]. These results refer to computations performed through quantum means that go beyond Turing's barrier. Calude and Pavlov show how a mathematical quantum device, which encodes the whole data into an infinite superposition, can be constructed to solve the Halting Problem [45]. Although their "halting machine" may not be of practical interest, the theoretical breakthrough is certainly important.

6.5 Provability

Long before quantum mechanics was suspected of enhancing the computational capabilities of classical Turing machines, Gödel showed that *truth* and *provability* are distinct concepts in any sufficiently strong formal system. Physics (and quantum mechanics in particular) makes no exception, since as a mathematical science it is treated formally. The introduction of quantum computers brings another interesting differentiation between the ability to prove and the ability to provide the proof trace. In principle, a QTM could be used to create some proof that relied upon quantum

mechanical interference among all the computational paths pursued in parallel in a superposition. We can measure the final state of the QTM in order to find an answer to our conjecture, but there is no way to record the intermediate steps of the proof without invariably disrupt its future course.

6.6 Complexity

Unlike *computability*, who tries to clearly distinguish between problems that computers can and cannot do, *complexity* focuses on the efficiency with which an answer to a solvable problem is computed. In other words, complexity quantifies how the memory and time resources scale with problem size. From the complexity point of view, **the most interesting open question is whether a QTM can make all NP problems tractable**, since many of these problems have very practical applications, but are computationally hard to solve.

Quantum Turing Machines can be thought of as quantum mechanical generalizations of probabilistic Turing machines (PTM). The key difference though is that in a PTM only one particular computational trajectory is followed (even if it's non-deterministic), while in a QTM all possible computational trajectories are pursued simultaneously, leading to a superposition of the achievable states and allowing complex phenomena, like quantum mechanical interference to take place. To support this behavior, each cell of the QTM's tape must be able to encode a blend of 0 and 1 simultaneously. Furthermore, the tape as a whole can exist in highly non-classical states, with different parts of it becoming entangled together as the computational process progresses. The head of the machine can also be in a *spatial* superposition state, because in order to be able to follow multiple computational paths simultaneously, the head must have the capacity to be at several locations at the same time.

These capabilities allowed the development of a new technique, which Deutsch called *quantum parallelism*, and which proved to be useful in computing some joint properties of all the outputs of a function faster than a classical Turing machine [61]. However, with respect to mere function calculation, Deutsch proved that QTMs have the same complexity class as TMs. Jozsa analyzed the power of quantum parallelism, giving a mathematical characterization for the classes of joint properties that can and cannot be computed by quantum parallelism [112].

6.6.1 QTM vs. DTM

Deutsch and Jozsa exhibited for the first time an example in which the QTM is exponentially faster than a deterministic Turing machine (DTM) [67]. Their quantum algorithm distinguishing between constant and balanced functions needs only one evaluation of the function, simultaneously on all possible inputs, while the classical machine has no other choice but to compute all function values sequentially and then counting the two possible outputs to see whether they balance. Despite its impressive speedup, the Deutsch-Jozsa algorithm is of no practical importance, having no known applications. Furthermore, the problem is also easy for a PTM, who can solve it very quickly with high probability. Therefore, the race was now on to find a problem for which a QTM could beat both a DTM and a PTM.

6.6.2 Quantum complexity classes

Deutsch and Jozsa were also the first to propose the use of quantum complexity classes to capture the difficulty of solving specific problems on quantum models of computation. Thus, in analogy with the classical classes P , ZPP (*Zero error Probability in Polynomial time*), and BPP (*Bounded error Probability in Polynomial time*) we have the quantum classes QP , ZQP and BQP . These

mean that a problem can be solved with certainty in worst-case polynomial time, with certainty in average-case polynomial time, and with probability greater than $2/3$ in worst-case polynomial time, respectively, by a quantum Turing machine. The work of Yao [185] who showed that complexity theory for quantum circuits matches that of QTMs legitimizes the study of quantum circuits, which are simpler to design and analyze than QTMs. Thus, the running time of a quantum algorithm is usually expressed as a function of the number of elementary operations, that is, elementary unitary transformations (quantum gates) that have to be applied to the initial state in order to evolve it into the final state from which the answer can be extracted through measurement.

6.6.3 QTM vs. PTM

The first hint that QTMs might be more powerful than PTMs was given by Bernstein and Vazirani, who showed how to sample from the Fourier spectrum of any Boolean function on n bits in polynomial time on a QTM [26]. No algorithm was known to replicate this result on a PTM. Then, Berthiaume and Brassard were able to construct an oracle, relative to which a decision problem exists that could be solved with certainty in polynomial time in the worst case on a quantum computer, but could not be solved classically in probabilistic expected polynomial time, if errors were not tolerated [28]. In the same paper, they also show that there is a decision problem solvable in exponential time on a QTM and in double exponential time on all but finitely many instances on any DTM. These two results, besides being a victory of quantum computers over classical machines (deterministic or probabilistic) also prove that the power of quantum computation cannot simply be ascribed to the indeterminism inherent in quantum theory.

Further evidence that QTMs are more powerful than PTMs was brought by Bernstein and Vazirani [26] who showed that there exist oracles under which there are problems belonging to BQP but not BPP . Moreover, Simon managed to prove the stronger result that there exists an oracle relative to which BQP cannot even be simulated by a PTM allowed to run for an exponential number of steps [165].

Unfortunately, all these results share the same drawbacks as Deutsch-Jozsa algorithm. In the first place, they are relativized results, so they do not break any major ground in terms of computational complexity. Note, in this context, that the quantum circuitry responsible for computing the function f in the Deutsch-Jozsa algorithm can also be assimilated with a black box or oracle. Secondly, they generally are contrived problems, defined with a specific theoretical purpose in mind and do not offer practical applications. The development of Shor's algorithms for factoring integers and computing discrete logarithms was important especially from this point of view. Although they were not the first quantum algorithms achieving an exponential speed-up relative to the corresponding best known classical algorithms, they sure have the potential to deliver a devastating blow to the security of currently used cryptographic codes. As for the first observation above, despite the fact that they do not rely on the existence of any oracle, they still fail to fulfill the dream of solving all NP problems efficiently. The reason is that **neither factoring nor computing discrete logarithms is known to be NP -complete, in spite of the general belief that they are not in P .**

6.6.4 Quantum versus classical complexity classes

The relative power of quantum computers with respect to classical ones can be couched in the relationships between classical and quantum complexity classes. The lack of a precise answer to the alleged superiority of quantum computation is also reflected in the difficulty to place quantum complexity classes among classical ones. Few such universal (unrelativized) results have been proven




Figure 10: Relationships between quantum and classical complexity classes.

so far. Naturally, the class BQP attracts most interest, being considered, even more so than QP , the class of all computational problems which can be solved efficiently on a quantum computer. Shor's algorithms, for instance, belong to BQP , while it is not known whether they are also in BPP or not. So, exactly where BQP fits with respect to P , BPP , NP and $PSPACE$ is as yet unknown. What is known is that $BPP \subseteq BQP$ (that is, quantum computers can efficiently solve all the problems that are tractable for a PTM) and $BQP \subseteq PSPACE$ (there are no problems outside of $PSPACE$ which quantum computers can solve efficiently) [26]. Consequently, from $P \subseteq BPP \subseteq BQP \subseteq PSPACE$ we can see that BQP lies somewhere between P and $PSPACE$ (see Figure 10). Thus, we know for sure that BQP contains all of P and BPP , **but whether it also contains some problems in $PSPACE$ that are not in NP , for example, remains an open question.**

The difficulty to settle these issues comes in part from unsolved classical complexity problems, **like the question whether $PSPACE$ is strictly bigger than P .** Note the important interdependence between quantum complexity issues and classical complexity theory. **If some day, quantum computers will be proved to be strictly more powerful than classical computers (that is, $BPP \subset BQP$) then it will follow that P is not equal to $PSPACE$.** Many researchers have unsuccessfully attempted to prove this latter result, which is an indication that it may be quite non-trivial to prove the strict superiority of quantum computers over classical ones, despite some evidence in favor of this proposition.

6.6.5 Quantum speed-up

The important change in the attitude of quantum complexity theorists relative to the speed-up gained by quantum computers when dealing with NP -complete problems is also of great relevance to the task of identifying those problems that are best suited for enabling a quick quantum solution. The discovery of fast quantum algorithms for factoring and computing discrete logarithms has raised great hopes for the development of an efficient way to tackle all NP problems through quantum

means. One approach pursued by some researchers exploits the special structure that NP -complete problems share, allowing complete solutions to be assembled out of smaller, partial solutions [183]. Grover, together with Cerf and Williams, employed the amplitude amplification technique from his unstructured quantum search algorithm to perform a search in a tree structure [46]. The algorithm works by nesting a quantum search in the middle of the tree with another in the fringe, in order to bias the search in the fringe of the tree in favor of only the extensions of the partial solutions. For an NP -complete problem that takes N^x steps classically, this quantum-mechanical approach will reduce the number of steps to roughly $\sqrt{N^x}$, where $x < 1$ depends on the degree of constrainedness of the problem. While still far from an exponential speed-up, this result may nevertheless be of practical interest.

Little by little, people realized that in many cases an exponential increase in efficiency, due to the use of quantum techniques of computation, is not possible. Beals et al. [10] proved bounds on the efficiency of quantum computers relative to classical deterministic computers for several computational problems. In many cases, if the quantum machine takes N steps, then the classical one takes at most $O(N^6)$ steps. Finally, Bennett, Bernstein, Brassard and Vazirani [16] have concluded that treating the circuitry relating to an NP -complete problem as an oracle and trying to use some form of quantum parallelism to search in parallel through all the possible solutions to the problem will not improve on the speed of Grover's algorithm. **Although this result does not prove that NP cannot be contained in BQP ,** it does establish the fact that there is no intrinsic property of quantum computation that will function like a black box to solve NP -complete problems. And since Grover's search algorithm was proved to be optimal, **the current belief is that the quadratic improvement may be the best we can get out of a quantum computer in these kinds of tasks** [158].

We should not forget, however, that Grover's algorithm is optimal for an unstructured search space and that deeper structures may exist, which can be exploited easier using the set of tools provided by quantum computation, to yield a polynomial running time solution. This observation can be made compatible with the conjectured quadratic tight upper bound on the speed-up provided by quantum computers on NP -complete problems, only if we assume that problems like factoring do not belong to this category. There are two interesting implications of this hypothesis. First, it should reinforce the justification for the efforts devoted to showing that factoring is tractable even through classical means of computation, for otherwise where would this problem stand? In the second place, **a question arises whether similar complex or well disguised structures exist in other problems, that prevented an efficient classical solution to be discovered so far, and yet clever quantum algorithms may be able to exploit them in order to produce a quick solution.**

7 Quantum error correction

As the field of quantum computation and quantum information was developing, a huge gap opened up between theory and experiment. While theorists were racing ahead to devise novel applications for idealized quantum machines endowed with thousands of gates, experiment has barely got beyond the stage of building a single quantum gate. The extreme fragility of quantum states makes it very difficult to maintain their coherence in order to perform useful computations.

A challenging task The task of running a quantum computer at a reasonable degree of accuracy is much more challenging than in the case of a digital computer. Digital information is much easier to protect against errors than the analog information implied by the continuous variables describing an

arbitrary quantum state. The laws of quantum mechanics prevent, in general, a direct application of the classical error-correction techniques. We cannot inspect (measure) at leisure the state of a quantum memory register to check whether an ongoing computation is not off track without the risk of altering the intended course of the computation. Moreover, because of the no-cloning theorem, quantum information cannot be amplified in the same way digital signals can. So it is no surprise that many researchers initially doubted the feasibility of quantum computers [125, 177]. Since then, research efforts have managed to demonstrate that dealing with errors in quantum states is possible, at least from the theoretical point of view, making from quantum error correction one of the most active research areas in quantum computation and quantum information.

7.1 Quantum errors

The decoherence problems responsible for the alteration of quantum states occur because keeping a quantum system isolated from its environment is virtually impossible. Therefore, over time, the quantum system tends to couple with the environment, causing quantum information to leak out. The simplest idea to avoid such a scenario is to perform the useful quantum computation before major external errors may occur. Errors that describe the result of a quantum system coupling to the environment are termed *external*, as opposed to *internal* errors, which arise when the input state is not prepared exactly as we intended or when the architecture of the quantum computer is not exactly correct. Fortunately, in the case of internal errors, Wojciech Zurek found that input errors do not grow with time and that architectural errors determine an error in the computation, which is proportional to the square of the size of the error [187]. Both of these types of internal errors grow much more slowly in a quantum computer than in a classical computer, so the real concern remains eluding the external errors.

The two processes mainly responsible for inducing external errors are *dissipation* and *decoherence*. A dissipative error process usually causes a qubit to flip by losing energy to its environment. The loss of coherence, on the other hand, is a much more subtle process. It affects the phase of a qubit, a pure quantum mechanical (that is, non-classical) feature, undetectable through direct measurement. The entanglement between the state of a quantum memory register and the state of its environment tends to randomize the relative phases of the possible states of the memory register. The immediate consequence is the annihilation of the interference effects, a powerful tool used in any genuine quantum computation. Decoherence⁴ is very problematic because of the speed with which it occurs, allowing little time for a useful quantum evolution before classical behavior effects take over. Joos has estimated the coherence time of different-sized systems under various environments [111]. His analysis shows that coherence time is most affected by temperature and interactions with surrounding gas particles. DiVincenzo also tried to give an estimation for the maximal number of computational steps that can be performed without losing coherence, but from the point of view of the possible “materials” used for the physical realization of a qubit [71].

Trying to choose the best suited materials for the components of a quantum computer that is to be operated at low temperatures and in vacuum represents the passive alternative in managing decoherence. Although this may be enough for certain tasks, a general-purpose quantum computer will necessarily require a more active approach in undoing whatever errors may arise during the computation. Correcting quantum errors certainly requires much more ingenuity than fixing classical bits, but the basic idea of using redundancy is still useful.

⁴The term “decoherence” is sometimes used to denote the alteration of a quantum state in general. Hopefully, the context in which it is used can always eliminate any possible ambiguity.

7.2 Error correction via symmetrization

The technique called *error correction via symmetrization* [29, 6] is yet another example of how the duality of quantum-mechanical laws can be exploited for the benefit of quantum computation. Although the measurement postulate severely restricts us in recycling techniques from classical error correction, it can still offer conceptually new ways of achieving error correction that are simply unavailable to classical computers. Error correction via symmetrization relies on the projective effect of measurements to do the job. The technique uses n quantum computers, each performing the same computation. Provided no errors occur, the joint state of the n computers is a symmetric one, lying somewhere in the small symmetric subspace of the entire possible Hilbert space. Devising a clever measurement that projects the joint state back into the symmetric subspace should be able to undo possible errors, without even knowing what the error is.

To achieve this, the n quantum computers need to be carefully entangled with a set of ancilla qubits placed in a superposition representing all possible permutations of n objects. In this way, the computation can be performed over all permutations of the computers simultaneously. Then, by measuring the ancilla qubits, the joint state of the n computers can be projected back into just the symmetric computational subspace, without the errors being measured explicitly. Peres has shown that this technique is most appropriate for correcting several qubits that are slightly wrong, rather than correcting a single qubit that is terribly wrong [150]. Error correction via symmetrization can be applied repeatedly, at regular time intervals, to avoid the accumulation of large errors and continually project the computation back into its symmetric subspace. Although the symmetrization algorithm is cumbersome and unattractive from the practical point of view, due to its low efficiency, it has nevertheless the important theoretical merit of proving that quantum error correction is, in principle, possible. Later on, it was even shown how the algorithm could be employed to improve frequency standards and the accuracy of atomic clocks [104].

7.3 Error-correcting codes

However, sudden large errors, like those induced by spontaneous emission, for example, require a different error-correction strategy based on the use of quantum error-correcting codes. Like in the classical case, the information contained in a qubit is spread out over several qubits so that damage to any one of them will not influence the outcome of the computation. In the quantum case, though, the encoding of the logical qubit is achieved through the use of specific resources, by entangling the logical qubit with several ancilla qubits. In this way, the information in the state of the qubit to be protected is spread among the correlations characterizing an entangled state. Paradoxically enough, entanglement with the environment can be fought back using quantum error-correcting codes based on entanglement [152].

Peter Shor's second major contribution to the advancement of quantum computation was the creation in 1995 of an algorithm that could correct any kind of error (amplitude and/or phase errors) affecting a single qubit in a 9-qubit code [162]. In a different approach, Steane studied the interference properties of multiple particle entangled states and managed to devise a shorter, 7-qubit code [168]. The number of qubits necessary for a perfect recovery from a single error was later squeezed down to a minimum of five [24, 124].

Naturally, in order to cope with more than one error at a time, it is necessary to use larger and more elaborate codes. The construction of all these codes is based on the surprising, yet beautiful idea of *digitizing the errors*. How can quantum errors be digitized when, as the variables they affect, they form a continuum? The answer lies in the linear nature of quantum mechanics. Any possible error affecting a single qubit can be expressed as a linear combination of no errors (I),

bit flip errors (X), phase errors (Z) and bit flip phase errors (Y). Generalizing to the case of a quantum register, an error can be written as $\sum_i e_i E_i$ for some error operators E_i and coefficients e_i . The error operators can be tensor products of the single-bit error transformations or more general multibit transformations. An error correcting code that can undo the effect of any error belonging to a set of correctable errors E_i will embed n data qubits (logical qubits) in $n + k$ code qubits (physical qubits). The joint state of the ensemble of code qubits is subject to an arbitrary error, mathematically expressed as a linear combination of the correctable error operators E_i .

To recover the original encoded state, a syndrome extraction operator has to be applied that uses some ancilla qubits to create a superposition of the error indices i corresponding to those correctable error operators E_i that have transformed the encoded state. Measuring only the ancilla qubits will collapse the superposition of errors, yielding only one index k . But because the ancilla qubits were entangled with the code qubits through the application of the syndrome extraction operator, the side effect of the measurement is that the corruption caused by all error transformations will be undone, save for the one corresponding to index k . Consequently, only one inverse error transformation is required in order to complete the recovery process. In essence, knowing how to deal with a set of fundamental error transformations allows us to tackle any linear combination of them by projecting it to one of the basis components. This process is referred to as *digitizing* or *discretizing* the errors. The book of Nielsen and Chuang [146] offers a detailed treatment of quantum codes, explaining how ideas from classical linear codes can be used to construct large classes of quantum codes, as the Calderbank-Shor-Steane (CSS) codes [43, 169], or the stabilizer codes (also known as additive quantum codes), which are even more general than the CSS codes and are based on the stabilizer formalism developed by Gottesman [91].

7.3.1 Scalability

The major drawback in using large and intricate quantum codes is that the corrective circuit itself is as much prone to errors as the quantum circuit responsible for the main computation. The more errors we are attempting to rectify, the more the complexity and length of the recovery procedure will increase (see [76] for some theoretical bounds on the relationship between the number of data qubits, the total number of entangled qubits and the maximal number of errors that can be tolerated). Thus, we can only increase the size of the error correction codes up to a certain cutoff point, past which no further gains in accuracy can be made.

7.3.2 Concatenated codes

One attempt to overcome this limitation are the *concatenated* codes. If a certain code uses n physical qubits to encode one logical qubit, a concatenated version of that code is obtained by further encoding each of the n qubits in another block of n . This hierarchical structure (tree) can be further expanded to accommodate as many levels as desired. An important theoretical result was proved for such concatenated codes. The *threshold theorem* states that by adding more levels of concatenation, the overall chance for an error can be made arbitrarily small, provided that the probability of an individual error is kept below a certain critical threshold [153]. What this result is saying is that, in principle, we can reliably perform an arbitrarily long quantum computation. The threshold level depends on many factors, like the specifics of the code used, the type of errors and whether errors occur more frequently in qubit storage or in gate processing. Because the process of error recovery within each tier of the hierarchy remains almost as simple as that of the original n -qubit code, there is only a small overhead in the size of the circuit necessary to ensure reliability.

Of course, the high cost of using concatenated codes lies in the exponential increase in the number of qubits with the number of levels added.

7.4 Fault-tolerance

Fault-tolerant quantum computation is another approach in trying to cope with the effects of noise both in the main computational circuit and the recovery circuitry. The operation of each quantum gate on the encoded data has to be carefully re-designed into a procedure for performing an encoded gate on the encoded state in such a way as to prevent error propagation and error accumulation. This will ensure that error correction will be effective at removing the errors. Since the fault-tolerant gates form a discrete set, part of the problem is to simulate any possible idealized quantum gates using only fault-tolerant procedures. For example, single-bit phase changes required for performing the Fourier transform in Shor's factoring algorithm can be approximated using combinations of fault-tolerant fixed-angle versions. Following this line of thought it is possible to perform a universal set of logical operations (the Hadamard, phase, controlled-NOT and $\pi/8$ gates are one choice, but there are others as well) using only fault-tolerant procedures.

The second aspect of fault-tolerance refers to the possibility of introducing errors on the encoded qubits by the error-corrective process itself. Peter Shor outlined a method of fault-tolerant recovery which uses extra ancilla qubits and some extra circuitry to double-check the diagnosis of errors [163]. On the other hand, John Preskill has identified five fault-tolerance criteria which, when met, ensure that failures during the procedure for error correction do not propagate to cause too many errors in the encoded data [153]. Fault-tolerance principles and techniques can be successfully combined with concatenation codes to make possible a more effective error correction in quantum computing and achieve the arbitrary level of accuracy guaranteed by the threshold theorem.

7.5 Topological quantum computing

The final possibility that we mention here to fight decoherence goes by the name of *topological quantum computing*. This method is the most speculative but potentially the most robust. Its aim is to achieve fault-tolerance at the very level of the physical processes responsible for the operation of a quantum gate. Topological quantum computing takes advantage of the non-local characteristics of some quantum interactions by encoding quantum information in a more global fashion. In this way, local disturbances will not affect the quality of the computation. A quantum interaction having the desired properties and which, therefore, could be used as a building block in the implementation of topological quantum gates is the Aharonov-Bohm effect [152]. It has even been showed that it is possible to construct universal quantum gates using only Aharonov-Bohm interactions [119].

The remarkable advances witnessed in quantum error correction have been acknowledged even by the most hardened skeptics and have transformed the prospects for making quantum computing a practical reality. However, **the question whether the theoretical results outlined in this section will eventually be brought to life by a physically realizable implementation is still to be decided by the years to come.**

8 Physical embodiments of a quantum computer

The first models for a quantum computer (Benioff's quantum upgrade of a classical reversible Turing machine [13], Feynman's quantum mechanical computer [81] and Deutsch's Universal Quantum

Turing Machine [61]) were abstract, theoretical tools designed to fathom the possibilities and limitations of quantum information processing, without any intention to give them a physical interpretation. An important step to move out of the "designer Hamiltonians" era towards a more practical approach was David Deutsch's introduction of quantum logic gates and circuits (or networks) [62]. Addressing the universality issue, researchers have discovered that single-qubit rotations supplemented with a two-qubit logic gate (such as controlled-NOT) are enough to perform any possible quantum computation [5]. But the surprising result related to this matter was the finding that almost any two-bit quantum gate is good enough to build a universal quantum computer [128, 65]. This *computation friendly* feature of quantum mechanics (the fact that the laws of physics support computational universality) was good news for experimentalists eager to build a quantum computer, because it meant that computation could be built out of almost any kind of interaction or physical process.

David DiVincenzo promulgates the following five criteria necessary for a practical quantum computer:

1. a scalable physical system with well-characterized qubits;
2. the ability to initialize the qubit state;
3. decoherence times much longer than the quantum gate operation time;
4. a universal set of quantum gates;
5. the ability to measure specific qubits.

It is worth noting that some powerful quantum calculations can be performed without any two-bit quantum gates at all. Two Carnegie-Mellon scientists showed that the quantum Fourier transform involved in Shor's factoring algorithm can be implemented using only single-bit quantum gates [93]. The qubit interactions normally required to carry on the computation were simulated in a semi-classical way: according to the results of measurements performed on certain qubits in the calculation, the phases of others are adjusted using single-bit gates. However, the modular exponentiation necessary to complete Shor's algorithm still requires the use of some two-bit quantum gates.

8.1 The quantum molecular switch

The first blueprint for a practical quantum computing machine was devised by Seth Lloyd [127], who was building upon the work of a group of German physicists regarding the realization of a molecular quantum switch [171]. Lloyd extended their model to include quantum superpositions, conditional logic and measurement techniques. His proposal was to use an array of weakly interacting quantum states as the basis for a quantum memory register. A concrete hardware platform for his design could be a heteropolymer. Each different atom in such a molecule could play the role of a qubit, with the ground state and a metastable excited state implementing the necessary binary logic.

The software for such a polymer machine would consist of sequences of laser light pulses carefully tuned on the resonance frequencies (the difference between the energy of the ground state $|0\rangle$ and the energy of the excited state $|1\rangle$) characterizing each different atom in the molecular chain. By varying the length of the pulse we can put the responding molecular unit into any kind of superposition of ground and excited states. In particular, if a laser pulse that flips the state of a certain qubit is called a π -pulse (because it rotates the state by 180°), then by applying a $\frac{\pi}{2}$ -pulse (a burst of light

that lasts only half the time of π -pulses) we will obtain an equally weighted superposition of $|0\rangle$ and $|1\rangle$. This is how a Hadamard gate or a square-root-of-NOT gate could be implemented.

Furthermore, Lloyd showed how the interatomic forces between 3 neighboring molecular units could detune the resonance frequency of the middle atom (or unit) sufficient enough to allow the implementation of a three-bit quantum gate. Finally, an extra, short-lived excited state could be used to measure the state of an atom, by detecting if a photon with a certain characteristic energy is emitted during the process. Similar schemes for manipulating qubits can be applied to virtually any system in which there are local interactions.

8.2 Ion traps

The *ion trap* scheme imagined by Cirac and Zoller [50] was considered at some point as the favorite runner in the quantum race. According to their design, a quantum memory register would be physically realized by using "fences" of electromagnetic fields to trap a number of ions within the central region of an evacuated chamber. Each imprisoned ion embodies a qubit, with the ground state representing $|0\rangle$ and a metastable state representing $|1\rangle$. Transitions between the two internal energy levels are obtained through optical means, by shining a pulse of light from a laser beam of the appropriate frequency onto the target ion. Lasers are also employed to cool down the ions into the ground state and thus, initialize them for the computation. The cooling system, known as *optical molasses*, can also be used to improve the accuracy of atomic clocks and create the conditions for novel states of matter known as Bose-Einstein condensates (see [42] pages 242–243).

Being charged particles, the ions are strongly coupled together by the combination of the electric repulsion between them and the squeezing effect induced by the electric fields of the trap. These antagonist forces create vibrational waves traveling among the trapped ions. Cirac and Zoller showed how the quantized collective motion of the ions inside the trap (their vibrations) can serve as a mechanism for implementing the necessary conditional logic. The ion trap system potentially offers enormous flexibility, allowing the qubits interacting in a quantum gate to be non-adjacent, due to the fact that the vibrations influence the whole system. The measurement process follows the procedure outlined by Lloyd for the polymer machine. The absence or presence of a fluoresced photon tells us which state the atom was in. The scheme requires an extra energy level that couples strongly to the lowest energy state.

The first quantum logic gate built on ion trap technology was produced in 1995 [141]. The experimenters managed to achieve a 90% success rate on the proper operation of a controlled-NOT gate. The original design of Cirac and Zoller was somewhat simplified, in that they used a single beryllium ion to encode both qubits. The ion's hyperfine internal energy levels (determined by the interaction between the spin of the single electron in the outermost shell and the ion's nucleus) were chosen to incarnate one qubit, while two vibrational energy levels of the ion as a whole were selected to represent the basis states of the second qubit. Thus, we can say that the one-bit-per-atom limit was surpassed in this case. The solution adopted avoided the technological problem of building small enough lasers to address each qubit individually. Unfortunately, this becomes a serious issue for any attempt to scale up the ion trap design. Also, the measured decoherence time safely allowed the completion of the controlled-NOT operation, but seemed insufficient for an extended computation.

The idea of using the electronic and motional states of a single ion (this time of the element Calcium) to encode a 2-qubit state was later exploited to implement the Deutsch-Jozsa algorithm on an ion-trap quantum computer [96]. And if scalable solutions to the ion-trap architecture will ever be found, then it will be possible, in principle, to factor a number with 385 bits using "just"

1926 ytterbium atoms and 30 billion laser pulses [107]. Therefore, **the current research effort is focused on developing architectures for a large-scale ion-trap quantum computer** [51, 116].

8.3 Flying qubits

Another proposal, which goes by the name of *flying qubit*-based quantum computers uses the cavity QED (quantum electrodynamics) technology to emulate the functionality of a controlled-NOT quantum gate. Quantum information is encoded in the polarization states of photons and, to facilitate their interaction, they are placed inside a small cavity with highly reflecting walls together with a drifting cesium atom. The spacing between the mirrors in the cavity can be adjusted to resonate with a particular transition between two energy levels of the cesium atom and the target and control photons. The control qubit is initially prepared in a circularly polarized state, either $|+\rangle$ (left or counterclockwise) or $|-\rangle$ (right or clockwise), while the target qubit is linearly polarized, meaning that it can be described by an equal superposition of a left and a right circularly polarized state, such as

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

It turns out that the $|+\rangle$ component of the target qubit is phase-shifted only if the control qubit has excited the cesium atom, which in turn, happens only when the control photons are circularly polarized in a parallel direction to the atom's spin. This conditional phase-shift is the basic building block to construct any quantum logic circuit, when supplemented with one-bit rotations.

Although quantum-phase gates based on cavity QED have been successfully realized experimentally [59, 176], it is a very challenging endeavor to extend this technology to complicated quantum circuits. An alternative would be to use beam splitters, phase shifters, single-photon sources and photo-detectors in an all-optical attempt to build a quantum computer [120, 148, 151]. Better still, the cavity QED and ion trap technologies could be combined to scale up the atom trap technology [52]. Since they are traveling very fast, photons could be used to transfer quantum information between distant trapped atoms (perhaps through fiber optics) with each of the multibit ion traps responsible for storing information and local processing. The cavity QED interactions would provide the necessary methods for exchanging quantum information between the two different carriers. The same goal could be achieved by using entanglement between a trapped atom and a photon [32].

8.4 NMR quantum computing

The most advanced quantum computer prototypes that seemed to have been built so far are based on *nuclear magnetic resonance* (NMR) spectroscopy, the same technique that is known in the medical setting as *magnetic resonance imaging* (MRI) [87, 56]. The method manipulates the nuclear spin states of the atoms in a molecule to represent the qubits. When the sample is subjected to a very powerful and uniform external magnetic field, the spins of the nuclei tend to occupy a lower energy state, where they point with the field. Electromagnetic pulses in the radio frequency range are employed to flip spins completely or at intermediate orientations. The energy levels associated with the spin states of a nucleus are slightly altered by the closest electron orbits of neighboring atoms (effect known as the *chemical shift*) and also by the nuclear magnetic field of neighboring nuclei (spin-spin coupling). Thus, the resonance peaks in the NMR spectrum of a particular type of nucleus may be shifted to reflect the chemical composition of the sample. Analyzing these subtle

differences in the spectra we may be able to identify the molecular structure of the sample. But for NMR quantum computing, chemical shifts offer a way to address each qubit in a molecule individually, while spin-spin coupling is the solution to achieve conditional logic.

What distinguishes an NMR quantum computer from the other approaches is its massive parallelism and ultra-high redundancy. Each molecule in the test tube can be considered as an individual processing unit with a local quantum memory register. Each "programming" radio frequency pulse is tuned to resonate with the same part in each molecule simultaneously. Thus, the quantum computation is carried out on a vast number of "processors", in parallel. It is true that the majority of them are affected by thermal noise and decoherence, but we don't care about this, as long as the NMR signals generated by random spin orientations average out to zero.

However, even after the removal of the random background, what we are left with is a thermal distribution of states rather than a pure quantum state. So, in order to properly initialize our computation, we are faced with the problem of selecting an ensemble of molecules that share the same quantum state. One possible way to achieve this is to sacrifice some nuclear spins on each molecule to check for quantum purity [87]. Alternatively, the *spatial averaging* method [56] does not expend precious qubits, but it is very sensitive to inaccuracies. To measure a qubit, we must first apply certain photon pulses and then interpret the absorption spectrum obtained by processing the signal collected by the coils surrounding the apparatus.

The NMR quantum computing technology was able to support the implementation of complex algorithms. Grover's search algorithm was run inside a chloroform sample in the form of a quantum database query experiment [48]. The Deutsch-Jozsa algorithm was successfully attempted by two independent groups [49, 110]. In 1999, the first implementation of the quantum Fourier transform was demonstrated. Although there were doubts cast on whether the NMR experiments were capable of producing truly entangled states [41], more people tried to produce working experimental realizations of increasingly complicated quantum computations. In December 2001, scientists at IBM's Almaden Research Center claimed to have performed the world's most complicated quantum-computer calculation at that time [179]. In their 7-qubit experiment, they controlled billions of custom-designed molecules to implement the simplest meaningful instance of Shor's algorithm for factoring the number 15. Considering the unprecedented control required over the seven spins during the calculation, this result is quite an achievement.

Unfortunately, as with the other designs for a practical quantum computer, huge obstacles have to be surpassed in order to **make NMR a scalable technology**. First, the size of the quantum memory register is restricted by the number of nuclear spins, and hence, atoms in a single molecule. Then, as more qubits are added to the system, the strength of the NMR signal decreases exponentially as the number of possible spin states for the whole molecule increases exponentially. In other words, the number of representative molecules in the initial, sought-after "pure" state (with all spins pointing with the field) decreases exponentially. Therefore, there is a trade-off between the computational power gained and the weakening of the output signal (which has to be maintained above noise level). In recent years, there have been a large number of experiments implementing various quantum algorithms on a small number of qubits [178, 130, 117, 73].

8.5 Solid state approaches

All proposed technologies we have discussed so far are more or less extravagant and represent a radical departure from the semiconductor technology on which today's computer industry is based. There are also some proposals that are challenging this point of view and try to upgrade the semiconductor technology to the quantum level.

The Australian physicist Bruce Kane designed a silicon-based quantum computer in which qubits are hosted by the nuclear spins of phosphorus atoms [115]. Each qubit can be addressed individually by applying a tiny voltage to a metal strip or *gate* placed above each phosphorus atom. The qubits can be made to interact indirectly via a coupling mediated by electron spins. Conditional interaction is achieved through another type of gate, controlling two adjacent nuclear spins. Measuring the spin state of a phosphorus nucleus involves the detection of a small current due to the electron migration between phosphorus ion donors, which in turn depends on the various states of the nuclear and electron spins and the strength of the gates.

Quantum dots In the same class of semiconductor devices that may one day become the key components of a quantum computer are the *quantum dots*. These are tiny islands of semiconducting material on a chip, typically measuring a few hundred atoms across and usually surrounded by an electrical insulator. By applying a negative voltage around these blobs it is possible to squeeze out some of the freely moving electrons, such that only one electron is left inside. The energy levels of this electron are discrete, so they can be used as representations for 0 and 1. Single-qubit rotations would be straightforward to implement using laser light pulses of various lengths tuned to the frequency corresponding to the energy difference between the ground state and the first excited state of the electron. Conditional logic could be achieved if two closely spaced quantum dots were subjected to an electric field. This will make the resonance frequency of one dot (the target qubit) sensitive to the state of the other dot [7]. Provided the resonance frequencies differ sufficiently, an appropriately tuned π -pulse could selectively flip the target qubit state, subject to the state of the control qubit, thus realizing a controlled-NOT quantum gate.

These solid state approaches are much more scalable with respect to the number of qubits, but they suffer more from decoherence. Consequently, they do not allow, at this stage, but some very simple quantum calculations to be completed before the loss of coherence, which is just another side of the scaling problem. Moreover, quantum dots and related devices need special manufacturing techniques and liquid helium temperatures to function properly. On the other hand, research into nanoscale technology may also have a huge impact on conventional silicon chip manufacturing, if the concept of single-electron transistors will reach the point of industry production. The colossal reductions in size and amount of power consumed by conventional memory and logic devices would allow the current trends in miniaturization to be sustained.

However, for those interested in seeing quantum dots becoming a serious candidate for doing real quantum computing experiments, **decoherence remains the main concern**. Daniel Loss and David DiVincenzo proposed a more robust way of exploiting quantum dots, by storing the quantum information in the spin orientation (rather than the energy level) of each electron in a quantum dot [131]. They showed how quantum tunneling, achieved by raising the voltage of a metal gate for a certain period of time, could be manipulated to simulate the functionality of a quantum XOR (equivalent to a controlled-NOT) gate. Stephen Hellberg [97] also addresses decoherence and the difficulty of generating local magnetic fields for single-qubit rotations.

Some hopes seem to be raised by the idea of using superconducting quantum dots, incorporating Josephson junctions to implement qubits [132]. An all silicon design for a quantum computer claims to take advantage of NMR successful aspects (ensemble measurement, radio frequency control, long decoherence times), but at the same time allows for more qubits and improved initialization [123]. A more recent proposal suggests spin-pair encoded qubits in silicon [166].

With the large number of research papers reporting new experimental realizations every month, it is difficult to predict which (if any) of the enumerated approaches will ever evolve into a viable

technology for building a practical and useful quantum computer. It may also happen that a novel, currently unknown technology will emerge, transforming quantum computing into a practical reality.

9 Final remarks

Since the early 1980's, when the field began to materialize, quantum computation and quantum information has come a long way to impose itself as a stand-alone multidisciplinary science. Indeed, today we can rightfully acknowledge the existence of a *quantum computer science* with deep ramifications into information theory, algorithms and complexity, cryptography, networking and communications, fault tolerance and error correction.

Quantum mechanics offers some really powerful tools that researchers have tried to exploit in order to make information processing more efficient and secure. Superpositions, quantum parallelism, interference and entanglement are ultimately responsible for the computational power attributed to a quantum mechanical machine. Similarly, the measurement principle and no-clonability theorem form the basis on which quantum cryptographic protocols were developed. Some quantum effects, especially entanglement and non-locality, seem to go so much against our common sense that even today there are voices who doubt their existence as "elements of reality" (to paraphrase Einstein) and have expressed reservations with respect to the experiments trying to "prove" them.

This ambiguous status of entanglement and the difficulty experimenters face in creating multibit entanglement prompted some researchers to analyze more rigorously the implications of entanglement for quantum computing. While proving that multi-partite entanglement is necessary for quantum algorithms operating on pure states in order to offer an exponential speed-up, Jozsa and Linden also argue that it is misleading to view entanglement as a key resource for quantum computational power [114]. Other researchers seem to have reached similar conclusions. Eli Biham, Gilles Brassard, Dan Kenigsberg and Tal Mor show that quantum computing without entanglement still provides some advantage over classical computing, but that, at the moment, entanglement is necessary for all practical purposes [30].

Counterfactual quantum computing However, seemingly absurd phenomena may occur even in the absence of entanglement. An example is *counterfactual* quantum mechanics. Counterfactuals, things that might have happened, although they did not in fact happen, have no physical consequences in classical physics. However, the potential occurrence of a quantum event can change the probabilities of obtaining certain experimental outcomes. Interaction-free measurements provide one possible context for quantum counterfactuals. An interaction-free measurement can be characterized as a kind of nondisturbing quantum measurement in which no energy is exchanged between the probe and the object [122]. Consequently, the result is rather inferred than obtained through direct inspection [68].

Elitzur and Vaidman [78] describe how a Mach-Zehnder interferometer can be used to perform an interaction-free measurement in the dramatic context of detecting an ultrasensitive bomb. The Mach-Zehnder interferometer (depicted in Figure 11) is an optical device composed of beam splitters, mirrors and photon detectors carefully placed to bring about quantum interference when a photon travels through the apparatus. Thus, when a photon enters the first beam splitter horizontally, it will always emerge from the horizontal port of the second beam splitter, provided the two arms of the interferometer have equal lengths. As in the case of Young's two-slit experiment, the reason is *self-interference*. Based on this phenomenon, Elitzur and Vaidman show that, if a bomb that will explode when hit by a single photon is placed in one arm of the interferometer, we




Figure 11: A Mach-Zehnder interferometer (BS=beam splitter; M=mirror).

still have a 25% chance of detecting it without causing it to explode. Furthermore, it is actually sufficient to couple the bomb to a sensor that performs a *quantum nondemolition* measurement on the photon, registering its passage but allowing it to continue to the final beam splitter. There are schemes through which the efficiency of this interaction-free measurement can be made arbitrarily close to one [121].

Richard Jozsa gave a computational spin to the interaction-free bomb detection thought experiment [113]. He replaced the bomb with a quantum computer capable of answering a certain decision problem, such as deciding whether a given number is prime, for example. The computer performs its computation when a photon is detected passing through that arm, otherwise it sits idle. Under these conditions, a quantum counterfactual result can be obtained again. In 25% of the cases when the tested number is prime, we can actually reach this conclusion without having to run the computer. So the mere fact that the computer is capable of producing the solution allows us to infer the correct answer without actually running the computer.

Exotic as it is, Deutsch's *many worlds interpretation* may be the most "intuitive" way to explain quantum counterfactuals. According to the multiverse interpretation, the computer did not run in our universe, but it did run in some other, parallel universe, therefore making the condition that the quantum computer has the potential to decide the question very important. We also note that there is a subtle quantum information exchange between the possible universes, which determines whether the interference phenomenon at the final beam splitter will take place or not. In his book "The Fabric of Reality" David Deutsch strongly advocates the "many worlds" interpretation, bringing a powerful argument. When a 250-digit number is factored using Shor's algorithm, the number of interfering universes is of the order of 10^{500} . Since there are only about 10^{80} atoms in the entire visible universe, he argues that "physical reality would not even remotely contain the resources required to factorize such a large number" ([63] page 217). As we have mentioned in the

opening section, the differences between the various interpretations of quantum mechanics are a matter of taste to physicists and ultimately involve philosophical notions like consciousness, which are beyond the scope of this paper.

Quantum mechanics is considered the most accurate description of the Universe we currently have, but in the light of possible future discoveries, we may have to adjust this theory some day. It is not clear now how the theoretical foundations of quantum computation and quantum information will be affected in such a case, but the optimistic view is that even in the worst case, the novel physical theory that will emerge may give rise to a new computational paradigm, maybe even more powerful than quantum computing.

Although from the computational complexity perspective, the question whether a quantum computer is ultimately more powerful than a classical machine was not given a clear and definitive answer yet (see the discussion in section 6), it would still make a big difference if a computing machine based on quantum principles could be built. The possibility of breaking today's public key cryptographic codes is certainly the most appealing, but even obtaining a quadratic speed-up for the hard NP problems would be a significant improvement of practical interest. The previous section has made us aware of the great challenges and obstacles towards building a practical quantum computer. Fortunately, in cryptography, the experiments designed to implement quantum protocols seem to be much more advanced. They could become the most visible touch of quantum technology in the short run.

But even if large-scale quantum computing machines will prove unfeasible, a quantum computer with only a relatively small number of qubits might still be of some use. Such a device could be the simulator Feynman first envisaged, a machine used to effectively emulate other quantum systems. Also, a small-scale quantum computer may eventually become a useful tool of experimental physics. The ability of creating and manipulating just a handful of qubits can allow physicists to run interesting tests on certain predictions of quantum theory, thus helping basic physics research at the very least.

Whatever the future of quantum information processing is, the most fundamental idea we can learn from it is that information is intrinsically physical and the power of a certain computational paradigm ultimately depends on the characteristics of the physical support used to embody information.

References

- [1] Daniel S. Abrams and Seth Lloyd. A quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. <http://xxx.lanl.gov/abs/quant-ph/9807070>, July 1998.
- [2] Amir D. Aczel. *Entanglement*. Raincoast Books, Vancouver, 2002.
- [3] Alain Aspect, J. Dalibard, and G. Roger. Experimental test of Bell's inequalities using time-varying analyzers. *Physical Review Letters*, 49:1804–1807, 1982.
- [4] Adriano Barenco. A universal two-bit gate for quantum computation. *Proceedings of the Royal Society of London A*, 449:679–683, 1995.
- [5] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, T. Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52:3457–3467, 1995. <http://arxiv.org/abs/quant-ph/9503016>.

- [6] Adriano Barenco, André Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. Stabilization of quantum computations by symmetrization. <http://xxx.lanl.gov/abs/quant-ph/9604028>, April 1996.
- [7] Adriano Barenco, David Deutsch, and Artur Ekert. Conditional quantum dynamics and logic gates. *Physical Review Letters*, 74(20):4083–4086, May 15, 1995. <http://arxiv.org/abs/quant-ph/9503017>.
- [8] Adriano Barenco and Artur Ekert. Quantum computation. *Acta Physica Slovaca*, 45:205–216, 1995.
- [9] Adriano Barenco, Artur Ekert, Kalle-Antti Suominen, and Päivi Törmä. Approximate Quantum Fourier Transform and Decoherence. *Physical Review A*, 54(1):139–146, July 1996.
- [10] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, pages 352–361, Palo Alto, CA, November 8–11, 1998. FOCS '98. <http://arxiv.org/abs/quant-ph/9802049>.
- [11] C. W. J. Beenakker and M. Kindermann. Quantum teleportation by particle-hole annihilation in the Fermi sea. *Physical Review Letters*, 92, 056801, February 6, 2004.
- [12] John Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.
- [13] Paul Benioff. The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines. *Journal of Statistical Physics*, 22:563–591, 1980.
- [14] Eric Benjamin, Kenny Huang, Amir Kamil, and Jimmy Kitiyachavalit. Quantum computability and complexity and the limits of quantum computation. <http://www.cs.berkeley.edu/kamil/quantum/qc4.pdf>, December 2003.
- [15] Charles H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6):525–532, 1973.
- [16] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh V. Vazirani. Strengths and weaknesses of quantum computing. *Special issue on Quantum Computation of the SIAM Journal on Computing*, 26(5):1510–1523, October 1997. <http://xxx.lanl.gov/abs/quant-ph/9701001>.
- [17] Charles H. Bennett, F. Bessette, and Gilles Brassard. Experimental quantum cryptography. In *Lecture Notes in Computer Science*, volume 473, pages 253–265. Springer-Verlag, Berlin, 1991.
- [18] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, IEEE, New York, 1984. Bangalore, India, December 1984.
- [19] Charles H. Bennett and Gilles Brassard. The dawn of a new era for quantum cryptography: The experimental prototype is working! *SIGACT News*, 20:78–82, Fall 1989.

- [20] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, A. Peres, and W. Wootters. Teleporting an unknown quantum state via dual classical and EPR channels. *Physical Review Letters*, 70:1895–1899, 1993.
- [21] Charles H. Bennett, Gilles Brassard, Sandu Popescu, Benjamin Schumacher, John A. Smolin, and William K. Wootters. Purification of noisy entanglement and faithful teleportation via noisy channels. *Physical Review Letters*, 76:722–725, 1996. <http://arxiv.org/abs/quant-ph/9511027>.
- [22] Charles H. Bennett, Gilles Brassard, and Jean-Marc Robert. Privacy amplification by public discussion. *SIAM Journal on Computing*, 17(2):210–229, April 1988.
- [23] Charles H. Bennett, David P. DiVincenzo, and John A. Smolin. Capacities of quantum erasure channels. *Physical Review Letters*, 78(16):3217–3220, 1997. <http://arxiv.org/abs/quant-ph/9701015>.
- [24] Charles H. Bennett, David P. DiVincenzo, John A. Smolin, and William K. Wootters. Mixed state entanglement and quantum error correction. *Physical Review A*, 54:3824–3851, 1996. <http://arxiv.org/abs/quant-ph/9604024>.
- [25] Charles H. Bennett and S. J. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physical Review Letters*, 69(20):2881–2884, 1992.
- [26] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *Special issue on Quantum Computation of the SIAM Journal on Computing*, 26(5):1411–1473, October 1997. <http://arxiv.org/abs/quant-ph/9701001>.
- [27] André Berthiaume. Quantum computation. In Lane A. Hemaspaandra and Alan L. Selman, editors, *Complexity Theory Retrospective II*, pages 23–51. Springer-Verlag, New York, 1997.
- [28] André Berthiaume and Gilles Brassard. Oracle quantum computing. *Journal of Modern Optics*, 41(12):2521–2535, December 1994.
- [29] André Berthiaume, David Deutsch, and Richard Jozsa. The stabilization of quantum computation. In *Proceedings of the Workshop on Physics and Computation: PhysComp '94*, pages 60–62, Los Alamitos, CA, 1994, 1994. IEEE Computer Society Press.
- [30] Eli Biham, Gilles Brassard, Dan Kenigsberg, and Tal Mor. Quantum computing without entanglement. <http://arxiv.org/abs/quant-ph/0306182>, June 2003.
- [31] David Biron et al. Generalized Grover search algorithm for arbitrary initial amplitude distribution. <http://xxx.lanl.gov/abs/quant-ph/9801066>, 1998.
- [32] B. B. Blinov, D. L. Moehring, L.-M. Duan, and Chris Monroe. Observation of entanglement between a single trapped atom and a single photon. *Nature*, 428:153–157, March 11, 2004.
- [33] Niels Bohr. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 48:696–702, October 1935.
- [34] D. Boschi, S. Branca, Francesco De Martini, L. Hardy, and Sandu Popescu. Experimental realization of teleporting an unknown pure quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 80:1121–1125, 1998. <http://arxiv.org/abs/quant-ph/9710013>.

- [35] D. Bouwmeester, J. W. Pan, Klaus Mattle, M. Eibl, Harald Weinfurter, and Anton Zeilinger. Experimental quantum teleportation. *Nature*, 390(6660):575–579, 1997.
- [36] Michel Boyer, Gilles Brassard, Peter Hoyer, and Alain Tapp. Tight bounds on quantum searching. In *Proceedings of the Workshop on Physics and Computation: PhysComp '96*, pages 36–43, Los Alamitos, CA, 1996, 1996. IEEE Computer Society Press. <http://xxx.lanl.gov/abs/quant-ph/9605034>.
- [37] Gilles Brassard and Claude Crépeau. 25 years of quantum cryptography. *SIGACT News*, 27(3):13–24, September 1996.
- [38] Gilles Brassard, Claude Crépeau, Richard Jozsa, and D. Langlois. A quantum bit commitment scheme provably unbreakable by both parties. In *Proceedings of the 34th Annual IEEE Symposium on Foundations of Computer Science*, pages 362–371. IEEE press, 1993.
- [39] Gilles Brassard, Claude Crépeau, and Stefan Wolf. Oblivious transfers and privacy amplification. *Journal of Cryptology*, 16(4):219–237, September 2003.
- [40] Gilles Brassard, Peter Hoyer, and Alain Tapp. Quantum counting. <http://xxx.lanl.gov/abs/quant-ph/9805082>, May 1998.
- [41] S. L. Braunstein, C. M. Caves, Richard Jozsa, N. Linden, Sandu Popescu, and R. Schack. Separability of very noisy mixed states and implications for NMR quantum computing. *Physical Review Letters*, 83:1054–1057, 1999. <http://xxx.lanl.gov/abs/quant-ph/9811018>.
- [42] Julian Brown. *The quest for the quantum computer*. Simon & Schuster, Touchstone edition, 2001.
- [43] A. R. Calderbank and Peter W. Shor. Good quantum error-correcting codes exist. *Physical Review A*, 54(2):1098–1106, 1996. <http://arxiv.org/abs/quant-ph/9512032>.
- [44] Cristian S. Calude, Michael J. Dinneen, and K. Svozil. Reflections on quantum computing. *Complexity*, 6:35–37, 2000.
- [45] Cristian S. Calude and Boris Pavlov. Coins, quantum measurements, and Turing’s barrier. *Quantum Information Processing*, 1(1–2):107–127, April 2002.
- [46] Nicolas J. Cerf, Lov K. Grover, and Colin P. Williams. Nested quantum search and NP-complete problems. *Physical Review A*, 61, 032303, 2000. <http://xxx.lanl.gov/abs/quant-ph/9806078>.
- [47] Andrew M. Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A. Spielman. Exponential algorithmic speedup by quantum walk. In *Proceedings of the 35th Annual ACM Symposium on the Theory of Computing*, pages 59–68, 2003. <http://arxiv.org/abs/quant-ph/0209131>.
- [48] Isaac L. Chuang, Neil Gershenfeld, and Mark Kubinec. Experimental implementation of fast quantum searching. *Physical Review Letters*, 80:3408–3412, 1998.
- [49] Isaac L. Chuang, Lieven M. K. Vandersypen, Xinlan Zhou, Debbie W. Leung, and Seth Lloyd. Experimental realization of a quantum algorithm. <http://xxx.lanl.gov/abs/quant-ph/9801037>, 1998.

- [50] Ignazio Cirac and Peter Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74:4091–4094, 1995.
- [51] Ignazio Cirac and Peter Zoller. A scalable quantum computer with ions in an array of microtraps. *Nature*, 404:579–581, April 6, 2000.
- [52] Ignazio Cirac, Peter Zoller, H. J. Kimble, and H. Mabuchi. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Physical Review Letters*, 78(16):3221–3224, April 21, 1997. <http://arxiv.org/abs/quant-ph/9611017>.
- [53] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 49:1804–1807, 1969.
- [54] Richard Cleve, Artur Ekert, Chiara Macchiavello, and Michele Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society of London A*, 454:339–354, 1998.
- [55] David G. Cory and Timothy F. Havel. Ion entanglement in quantum information processing. *Science*, 304(5676):1456–1457, June 4, 2004.
- [56] David G. Cory, Mark D. Price, and Timothy F. Havel. Nuclear magnetic resonance spectroscopy: An experimentally accessible paradigm for quantum computing. <http://xxx.lanl.gov/abs/quant-ph/9709001>, 1997.
- [57] Claude Crépeau. Efficient cryptographic protocols based on noisy channels. *Lecture Notes in Computer Science*, 1233:306–317, 1997.
- [58] Claude Crépeau and Joe Kilian. Achieving oblivious transfer using weakened security assumptions. In *Proceedings of the 29th Annual IEEE Symposium on Foundations of Computer Science*, pages 42–52. IEEE press, October 1988.
- [59] L. Davidovich et al. Quantum switches and nonlocal microwave fields. *Physical Review Letters*, 71(15):2360–2363, October 11, 1993.
- [60] David Deutsch. Quantum theory as a universal physical theory. *International Journal of Theoretical Physics*, 24(1):1–41, 1985.
- [61] David Deutsch. Quantum theory, the Church-Turing principle, and the Universal Quantum Computer. *Proceedings of the Royal Society of London A*, 400:97–117, 1985.
- [62] David Deutsch. Quantum computational networks. *Proceedings of the Royal Society of London A*, 425:73–90, 1989.
- [63] David Deutsch. *The fabric of reality*. Penguin Books, 1998.
- [64] David Deutsch. The structure of the multiverse. <http://xxx.lanl.gov/abs/quant-ph/0104033>, April 2001.
- [65] David Deutsch, Adriano Barenco, and Artur Ekert. Universality in quantum computation. *Proceedings of the Royal Society of London A*, 449(1937):669–677, June 8, 1995.
- [66] David Deutsch, Artur Ekert, Richard Jozsa, Chiara Macchiavello, Sandu Popescu, and Anna Sanpera. Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Physical Review Letters*, 77:2818–2821, 1996. <http://arxiv.org/abs/quant-ph/9604039>.

- [67] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London A*, 439:553–558, 1992.
- [68] Robert Dicke. Interaction-free quantum measurements: A paradox? *American Journal of Physics*, 49:925–930, November 1981.
- [69] Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, November 1976.
- [70] P. Dirac. *The Principles of Quantum Mechanics*. Oxford University Press, 4th edition, 1958.
- [71] David DiVincenzo. Quantum computation. *Science*, 270:255–261, October 13, 1995.
- [72] David DiVincenzo. Two-bit gates are universal for quantum computation. *Physical Review A*, 51:1015–1022, 1995.
- [73] Jiangfeng Du et al. Experimental implementation of the quantum random-walk algorithm. *Physical Review A*, 67, 042316, 2003.
- [74] Albert Einstein, Boris Podolsky, and Nathan Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 47:777–780, May 1935.
- [75] Artur Ekert. Quantum cryptography based on Bell’s theorem. *Physical Review Letters*, 67:661–663, 1991.
- [76] Artur Ekert and Chiara Macchiavello. Quantum error correction for communication. *Physical Review Letters*, 77:2585–2588, 1996.
- [77] Artur Ekert, John Rarity, P. Tapster, and G. Palma. Practical quantum cryptography based on two-photon interferometry. *Physical Review Letters*, 69:1293–1295, 1992.
- [78] Avshalom Elitzur and Lev Vaidman. Quantum-mechanical interaction-free measurements. *Foundations of Physics*, 23:987–997, 1993.
- [79] Edward Farhi and Sam Gutmann. Quantum computation and decision trees. *Physical Review A*, 58(2):915–928, August 1998. <http://arxiv.org/abs/quant-ph/9706062>.
- [80] Richard Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6 & 7):467–488, 1982.
- [81] Richard Feynman. Quantum mechanical computers. *Foundations of Physics*, 16(6):507–531, 1986.
- [82] Richard Feynman, R. B. Leighton, and M. Sands. *The Feynman Lectures on Physics*, volume III. Addison-Wesley, Reading, Mass., 1965.
- [83] Lance Fortnow. One complexity theorist’s view of quantum computing. *Theoretical Computer Science*, 292(3):597–610, 2003.
- [84] J. Franson and H. Ilves. Quantum cryptography using optical fibers. *Applied Optics*, 33:2949–2954, 1995.
- [85] Ed Fredkin and Tom Toffoli. Conservative logic. *International Journal of Theoretical Physics*, 21(3/4):219–253, 1982.

- [86] A. Furusawa, J. L. Sørensen, S. L. Braunstein, C. A. Fuchs, H. J. Kimble, and E. S. Polzik. Unconditional quantum teleportation. *Science*, 282:706–709, 1998.
- [87] Neil Gershenfeld and Isaac L. Chuang. Quantum computing with molecules. *Scientific American*, pages 66–71, June 1998.
- [88] Nicolas Gisin, Renato Renner, and Stefan Wolf. Linking classical and quantum key agreement: Is there a classical analog to bound entanglement? *Algorithmica*, 34(4):389–412, 2002.
- [89] Nicolas Gisin and Stefan Wolf. Quantum cryptography on noisy channels: quantum versus classical key-agreement protocols. *Physical Review Letters*, 83:4200–4203, 1999. <http://arxiv.org/abs/quant-ph/9902048>.
- [90] Nicolas Gisin and Stefan Wolf. Linking classical and quantum key agreement: Is there "bound information"? <http://arxiv.org/abs/quant-ph/0005042>, May 2000.
- [91] Daniel Gottesman. Class of quantum error-correcting codes saturating the quantum hamming bound. *Physical Review A*, 54:1862–1868, 1996. <http://arxiv.org/abs/quant-ph/9604038>.
- [92] Daniel M. Greenberger, Michael A. Horne, Abner Shimony, and Anton Zeilinger. Bell's theorem without inequalities. *American Journal of Physics*, 58(12):1131–1143, December 1990.
- [93] Robert Griffiths and Chi-Sheng Niu. Semiclassical fourier transform for quantum computation. *Physical Review Letters*, 76:3228–3231, 1996.
- [94] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing*, pages 212–219, Philadelphia, Pennsylvania, 22–24 May 1996, 1996.
- [95] Lov K. Grover. A framework for fast quantum mechanical algorithms. In *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing*, pages 53–62, May 1998. <http://xxx.lanl.gov/abs/quant-ph/9711043>.
- [96] Stephan Gulde et al. Implementation of the Deutsch-Jozsa algorithm on an ion-trap quantum computer. *Nature*, 421:48–50, January 2, 2003.
- [97] C. Stephen Hellberg. Robust quantum computation with quantum dots. <http://arxiv.org/abs/quant-ph/0304150>, April 23, 2003.
- [98] Mika Hirvensalo. *Quantum Computing*. Springer-Verlag, 2001.
- [99] Tad Hogg. Quantum computing and phase transitions in combinatorial search. *Journal of Artificial Intelligence Research*, 4:91–128, 1996. <http://xxx.lanl.gov/abs/quant-ph/9508012>.
- [100] Tad Hogg. Highly structured searches with quantum computers. *Physical Review Letters*, 80(11):2473–2476, March 1998.
- [101] Tad Hogg. Quantum search heuristics. *Physical Review A*, 61, 052311, 2000.
- [102] Tad Hogg and Dmitriy Portnov. Quantum optimization. *Information Sciences*, 128:181–197, 2000.

- [103] A. S. Holevo. The capacity of the quantum channel with general signal states. *IEEE Transactions on Information Theory*, 44(1):269–273, 1998.
- [104] Susanna F. Huelga, Chiara Macchiavello, Thomas Pellizzari, Artur K. Ekert, M. B. Plenio, and J. I. Cirac. Improvement of frequency standards with quantum entanglement. *Physical Review Letters*, 79(20):3865–3868, November 17, 1997.
- [105] Richard Hughes et al. Quantum cryptography. *Contemporary Physics*, 36(3):149–163, 1995. <http://arxiv.org/abs/quant-ph/9504002>.
- [106] Richard Hughes et al. Secure communications using quantum cryptography. *Proceedings of SPIE*, 3076:2–11, 1997.
- [107] Richard J. Hughes et al. Decoherence bounds on quantum computation with trapped ions. <http://xxx.lanl.gov/abs/quant-ph/9604026>, April 1996.
- [108] Richard J. Hughes, George L. Morgan, and C. Glen Peterson. Quantum key distribution over a 48 km optical fibre network. *Journal of Modern Optics*, 47(2/3):533–547, 2000.
- [109] Richard J. Hughes, Jane E. Nordholt, Derek Derkacs, and Charles G. Peterson. Practical free-space quantum key distribution over 10 km in daylight and at night. *New Journal of Physics*, 4:43.1–43.14, 2002.
- [110] Jonathan A. Jones and Michele Mosca. Implementation of a quantum algorithm to solve Deutsch’s problem on a nuclear magnetic resonance quantum computer. *Journal of Chemical Physics*, 109:1648–1653, 1998. <http://xxx.lanl.gov/abs/quant-ph/9801027>.
- [111] E. Joos and H. D. Zeh. The emergence of classical properties through interaction with the environment. *Zeitschrift für Physik B*, 59:223–243, 1985.
- [112] Richard Jozsa. Characterizing classes of functions computable by quantum parallelism. *Proceedings of the Royal Society of London A*, 435:563–574, 1991.
- [113] Richard Jozsa. Quantum effects in algorithms. <http://xxx.lanl.gov/abs/quant-ph/9805086>, May 29, 1998.
- [114] Richard Jozsa and Noah Linden. On the role of entanglement in quantum computational speed-up. <http://arxiv.org/abs/quant-ph/0201143>, 2002.
- [115] Bruce E. Kane. A silicon-based nuclear spin quantum computer. *Nature*, 393:133–137, May 14, 1998.
- [116] D. Kielpinski, Chris Monroe, and David J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709–711, June 13, 2002.
- [117] Jachyun Kim, Jae-Seung Lee, and Soonchil Lee. Experimental realization of a target-accepting quantum search by NMR. *Physical Review A*, 65, 054301, April 24, 2002.
- [118] A. Y. Kitaev. Quantum measurements and the Abelian stabilizer problem. <http://arxiv.org/abs/quant-ph/9511026>, November 1995.
- [119] A. Y. Kitaev. Fault-tolerant quantum computation by anyons. <http://xxx.lanl.gov/abs/quant-ph/9707021>, July 1997.

- [120] Emanuel H. Knill, Raymond Laflamme, and Gerard J. Milburn. A scheme for efficient quantum computation with linear optics. *Nature*, 409:46–52, January 4, 2001.
- [121] Paul G. Kwiat et al. High-efficiency quantum interrogation measurements via the quantum Zeno effect. *Physical Review Letters*, 83(23):4725–4728, December 6, 1999.
- [122] Paul G. Kwiat, Harald Weinfurter, T. Herzog, and Anton Zeilinger. Interaction-free measurement. *Physical Review Letters*, 74:4763–4766, 1995.
- [123] Thaddeus D. Ladd et al. An all silicon quantum computer. <http://arxiv.org/abs/quant-ph/0109039>, September 7, 2001.
- [124] Raymond Laflamme, Cesar Miquel, Juan Pablo Paz, and Wojciech Hubert Zurek. Perfect quantum error correction code. <http://arxiv.org/abs/quant-ph/9602019>, February 1996.
- [125] Rolf Landauer. Is quantum mechanics useful? *Philosophical Transactions of the Royal Society of London. Series A*, 353(1703):367–376, 1995.
- [126] Arjen K. Lenstra and H. W. Lenstra, Jr., editors. *The Development of the Number Field Sieve*. Springer-Verlag, New York, 1993.
- [127] Seth Lloyd. A potentially realizable quantum computer. *Science*, 261:1569–1571, September 17, 1993.
- [128] Seth Lloyd. Almost any quantum logic gate is universal. *Physical Review Letters*, 75(2):346–349, July 10, 1995.
- [129] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, August 1996.
- [130] Gui Lu Long and Li Xiao. Experimental realization of a fetching algorithm in a 7 qubit NMR quantum computer. <http://arxiv.org/abs/quant-ph/0207079>, July 15, 2002.
- [131] Daniel Loss and David P. DiVincenzo. Quantum computation with quantum dots. *Physical Review A*, 57(1):120–126, 1998. <http://arxiv.org/abs/cond-mat/9701055>.
- [132] Yuriy Makhlin, Gerd Scöhn, and Alexander Shnirman. Josephson-junction qubits with controlled couplings. *Nature*, 396:305–307, March 25, 1999.
- [133] Christophe Marand and Paul Townsend. Quantum key distribution over distances as long as 30 km. *Optics Letters*, 20(16):1695–1697, August 1995.
- [134] T. W. Marshall, E. Santos, and F. Selleri. Local realism has not been refuted by atomic cascade experiments. *Physics Letters A*, 98:5–9, 1983.
- [135] Klaus Mattle, Harald Weinfurter, Paul G. Kwiat, and Anton Zeilinger. Dense coding in experimental quantum communication. *Physical Review Letters*, 76(25):4656–4659, 1996.
- [136] Dominic Mayers. The trouble with quantum bit commitment. <http://xxx.lanl.gov/abs/quant-ph/9603015>, March 1996.
- [137] Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters*, 78:3414–3417, April 1997.

- [138] Ralph Merkle. Secure communications over insecure channels. *Communications of the ACM*, 21:294–299, 1978.
- [139] N. David Mermin. From Cbits to Qbits: Teaching computer scientists quantum mechanics. <http://arxiv.org/abs/quant-ph/0207118>, July 2002.
- [140] Gary L. Miller. Riemann’s hypothesis and tests for primality. *Journal of Computer and System Sciences*, 13(3):300–317, December 1976.
- [141] Chris Monroe et al. Demonstration of a fundamental quantum logic gate. *Physical Review Letters*, 75(25):4714–4717, 1995.
- [142] A. Muller, H. Zbinden, and N. Gisin. Quantum cryptography over 23 km in installed under-lake telecom fibre. *Europhysics Letters*, 33:335–339, 1996.
- [143] Marius Nagy and Selim G. Akl. On the importance of parallelism for quantum computation and the concept of a universal computer. Technical Report 2005–495, School of Computing, Queen’s University, Kingston, Ontario, May 2005. 18 pages.
- [144] Richard J. Nelson, David G. Cory, and Seth Lloyd. Experimental demonstration of Greenberger-Horne-Zeilinger correlations using nuclear magnetic resonance. *Physical Review A*, 61, 022106, 2000. 5 pages.
- [145] Michael A. Nielsen. Simple rules for a complex quantum world. *Scientific American*, 13(1):24–33, 2003. Special edition: The edge of physics.
- [146] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [147] Michael A. Nielsen, Emanuel H. Knill, and Raymond Laflamme. Complete quantum teleportation using nuclear magnetic resonance. *Nature*, 396(6706):52–55, 1998. <http://xxx.lanl.gov/abs/quant-ph/9811020>.
- [148] J. L. O’Brien et al. Demonstration of an all-optical quantum controlled-NOT gate. *Nature*, 426:264–267, November 20, 2003.
- [149] B. Pablo-Norman and M. Ruiz-Altaba. Noise in grover’s quantum search algorithm. *Physical Review A*, 61, 012301, January 2000. 5 pages.
- [150] Asher Peres. Error symmetrization in quantum computers. <http://xxx.lanl.gov/abs/quant-ph/9605009>, May 1996.
- [151] T. B. Pittman, M. J. Fitch, B. C. Jacobs, and J. D. Franson. Experimental controlled-NOT logic gate for single photons in the coincidence basis. *Physical Review A*, 68, 032316, 2003.
- [152] John Preskill. Fault-tolerant quantum computation. In Hoi-Kwong Lo, Sandu Popescu, and Tim Spiller, editors, *Introduction to quantum computation and information*, pages 213–269. World Scientific, 1998. <http://xxx.lanl.gov/abs/quant-ph/9712048>.
- [153] John Preskill. Reliable quantum computers. *Proceedings of the Royal Society of London A*, 454:385–410, 1998. <http://xxx.lanl.gov/abs/quant-ph/9705031>.

- [154] Michael O. Rabin. How to exchange secrets by oblivious transfer. Technical Report TR-81, Aiken Computation Laboratory, Harvard University, 1981.
- [155] Eleanor Rieffel and Wolfgang Polak. An introduction to quantum computing for non-physicists. *ACM Computing Surveys*, 32(3):300–335, September 2000.
- [156] Ronald L. Rivest, Adi Shamir, and Len M. Adleman. A method of obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.
- [157] H. P. Robertson. The uncertainty principle. *Physical Review*, 34:163–164, 1929.
- [158] Sara Robinson. Emerging insights on limitations of quantum computing shape quest for fast algorithms. *SIAM News*, 36(1), January/February 2003.
- [159] B. Schumacher. Quantum coding. *Physical Review A*, 51:2738–2747, 1995.
- [160] B. Schumacher and M. D. Westmoreland. Sending classical information via noisy quantum channels. *Physical Review A*, 56(1):131–138, 1997.
- [161] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana, 1949.
- [162] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52:2493–2496, October 1995.
- [163] Peter W. Shor. Fault-tolerant quantum computation. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 56–65, Los Alamitos, CA, 1996, 1996. IEEE Computer Society Press.
- [164] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *Special issue on Quantum Computation of the SIAM Journal on Computing*, 26(5):1484–1509, October 1997.
- [165] Dan R. Simon. On the power of quantum computation. *Special issue on Quantum Computation of the SIAM Journal on Computing*, 26(5):1474–1483, October 1997.
- [166] A. J. Skinner, M. E. Davenport, and Bruce E. Kane. Hydrogenic spin quantum computing in silicon: A digital approach. *Physical Review Letters*, 90, 087901, 2003.
- [167] Warren D. Smith. The energy-time uncertainty principle. <http://external.nj.nec.com/homepages/wds/etuncert2TR.ps>, typeset 811 September 21 1999.
- [168] Andrew M. Steane. Error correcting codes in quantum theory. *Physical Review Letters*, 77(5):793–797, July 29, 1996.
- [169] Andrew M. Steane. Multiple particle interference and quantum error correction. *Proceedings of the Royal Society of London A*, 452:2551–2576, 1996.
- [170] P. R. Tapster, J. G. Rarity, and P. C. M. Owens. Violation of Bell’s inequality over 4 km of optical fibre. *Physical Review Letters*, 73:1923–1926, 1994.
- [171] W. Teich, K. Obermayer, and G. Mahler. Structural basis for multistationary quantum systems ii: Effective few-particle dynamics. *Physical Review B*, 37(14):8111–8120, 1988.

- [172] Caroline H. Thompson. The Chaotic Ball: An intuitive model for EPR experiments. <http://xxx.lanl.gov/abs/quant-ph/9611037>, November 1996.
- [173] Caroline H. Thompson. Timing, "accidentals" and other artifacts in EPR experiments. <http://xxx.lanl.gov/abs/quant-ph/9711044>, November 1997.
- [174] Caroline H. Thompson. The tangled methods of quantum entanglement experiments. *Accountability in Research*, 6(4):311–332, 1999.
- [175] W. Tittel, J. Brendel, B. Gisin, T. Herzog, H. Zbinden, and N. Gisin. Experimental demonstration of quantum-correlations over more than 10 kilometers. *Physical Review A*, 57:3229–3232, 1998. <http://xxx.lanl.gov/abs/quant-ph/9707042>.
- [176] Q. Turchette et al. Measurement of conditional phase shifts for quantum logic. *Physical Review Letters*, 75(25):4710–4713, 1995.
- [177] William G. Unruh. Maintaining coherence in quantum computers. *Physical Review A*, 51:992–997, 1995. <http://arxiv.org/abs/hep-th/9406058>.
- [178] Lieven M. K. Vandersypen, Matthias Steffen, Gregory Breyta, Constantino S. Yannoni, Richard Cleve, and Isaac L. Chuang. Experimental realization of an order-finding algorithm with an NMR quantum computer. *Physical Review Letters*, 85(25):5452–5455, December 18, 2000.
- [179] Lieven M. K. Vandersypen, Matthias Steffen, Gregory Breyta, Constantino S. Yannoni, Mark H. Sherwood, and Isaac L. Chuang. Experimental realization of Shor’s quantum factoring algorithm using nuclear magnetic resonance. *Nature*, 414:829–938, December 2001.
- [180] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15:78–88, 1983.
- [181] Colin P. Williams and Scott H. Clearwater. *Explorations in quantum computing*. Springer-Verlag, New York, 1998.
- [182] Colin P. Williams and Scott H. Clearwater. *Ultimate zero and one: computing at the quantum frontier*. Springer-Verlag, New York, 2000.
- [183] Colin P. Williams and Tad Hogg. Exploiting the deep structure of constraint problems. *Artificial Intelligence Journal*, 70:73–117, 1994.
- [184] William K. Wootters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.
- [185] Andrew Chi-Chih Yao. Quantum circuit complexity. In *Proceedings of the 34th Annual Symposium on Foundations of Computer Science*, pages 352–360, Los Alamitos, CA, 1993, 1993. IEEE Computer Society Press.
- [186] Christof Zalka. Grover’s quantum searching algorithm is optimal. *Physical Review A*, 60:2746–2751, 1999. <http://xxx.lanl.gov/abs/quant-ph/9711070>.
- [187] Wojciech H. Zurek. Reversibility and stability of information processing systems. *Physical Review Letters*, 53:391–394, 1984.